

SEQUENTIAL TRACKING OF A HIDDEN MARKOV CHAIN USING POINT PROCESS OBSERVATIONS

ERHAN BAYRAKTAR AND MICHAEL LUDKOVSKI

ABSTRACT. We study finite horizon optimal switching problems for hidden Markov chain models with point process observations. The controller possesses a finite range of strategies and attempts to track the state of the unobserved state variable using Bayesian updates over the discrete observations. Such a model has applications in economic policy making, staffing under variable demand levels and generalized Poisson disorder problems. We show regularity of the value function and explicitly characterize an optimal strategy. We also provide an efficient numerical scheme and illustrate our results with several computational examples.

1. INTRODUCTION

An economic agent (henceforth the controller) observes a compound Poisson process X with arrival rate λ , and mark/jump distribution ν . The local characteristics (λ, ν) of X are determined by the current state of an *unobservable* Markov jump process M with finite state space $E \triangleq \{1, \dots, m\}$. More precisely, the characteristics are (λ_i, ν_i) whenever M is at state i , for $i \in E$.

The objective of the controller is to *track* the state of M given the information in X . To do so, the controller possesses a range of policies a in the finite alphabet $\mathcal{A} \triangleq \{1, \dots, A\}$. The policies are sequentially adopted starting from time 0 and until some fixed horizon $T < \infty$. The infinite horizon case $T = +\infty$ is treated in Section 5.1. The selected policy a leads to running costs (benefits) at instantaneous rate

$$\sum_{i \in E} c_i(a) 1_{\{M_t=i\}} dt.$$

The controller's overall strategy consists of a double sequence (τ_k, ξ_k) , $k = 0, 1, 2, \dots$, with $\xi_k \in \mathcal{A}$ representing the sequence of chosen policies and $0 \triangleq \tau_0 < \tau_1 < \dots \leq T$ representing the times of policy changes (from now on termed *switching times*). We denote the entire strategy by the right-continuous piecewise constant process $\xi: [0, T] \times \Omega \rightarrow \mathcal{A}$, with $\xi_t = \xi_k$ if $\tau_k \leq t < \tau_{k+1}$ or

$$(1.1) \quad \xi_t = \sum_{\tau_{k+1} \leq T} \xi_k \cdot 1_{[\tau_k, \tau_{k+1})}(t).$$

2000 *Mathematics Subject Classification.* Primary 62L10; Secondary 62L15, 62C10, 60G40.

Key words and phrases. Markov modulated Poisson processes, optimal switching.

E. Bayraktar is supported in part by the National Science Foundation, under grant DMS-0604491.

Beyond running benefits, the controller also faces switching costs in changing her policy which lead to inertia and hysteresis. If at time t , the controller changes her policy from a to b and $M_t = i$ then an immediate cost $K_i(a, b)$ is incurred. The overall objective of the controller is to maximize the total present value of all tracking benefits minus the switching costs which is given by

$$\int_0^T e^{-\rho t} \left(\sum_{i \in E} c_i(\xi_t) 1_{\{M_t=i\}} \right) dt - \sum_k e^{-\rho \tau_k} \left(\sum_{i \in E} K_i(\xi_{\tau_k-}, \xi_{\tau_k}) \cdot 1_{\{M_{\tau_k}=i\}} \right),$$

where $\rho \geq 0$ is the discount factor.

Since M is unobserved, the controller must carry out a filtering procedure. We postulate that she collects information about M via a Bayesian framework. Let $\vec{\pi} = (\pi_1, \dots, \pi_m) \triangleq (\mathbb{P}\{M_0 = 1\}, \dots, \mathbb{P}\{M_0 = m\})$ be the initial (prior) beliefs of the controller about M and $\mathbb{P}^{\vec{\pi}}$ the corresponding conditional probability law. The controller starts with beliefs π , observes X , updates her beliefs and adjusts her policy accordingly. Because only X is observable, the strategy ξ should be determined by the information generated by X , namely each τ_k must be a stopping time of the filtration \mathcal{F}^X of X . Similarly, the value of each ξ_k is determined by the information $\mathcal{F}_{\tau_k}^X$ revealed by X until τ_k . These notions and the precise updating mechanism will be formalized in Section 2.3. We denote by $\mathcal{U}(T)$ the set of all such *admissible strategies* on a time interval $[0, T]$. Since strategies with infinitely many switches would have infinite costs, we exclude them from $\mathcal{U}(T)$.

Starting with initial policy $a \in \mathcal{A}$ and beliefs $\vec{\pi}$, the performance of a given policy $\xi \in \mathcal{U}(T)$ is

$$(1.2) \quad J^\xi(T, \vec{\pi}, a) \triangleq \mathbb{E}^{\vec{\pi}, a} \left[\int_0^T e^{-\rho t} \left(\sum_{i \in E} c_i(\xi_t) 1_{\{M_t=i\}} \right) dt - \sum_k e^{-\rho \tau_k} \left(\sum_{i \in E} K_i(\xi_{\tau_k-}, \xi_{\tau_k}) \cdot 1_{\{M_{\tau_k}=i\}} \right) \right].$$

The first argument in J^ξ is the remaining time to maturity. The optimization problem is to compute

$$(1.3) \quad U(T, \vec{\pi}, a) \triangleq \sup_{\xi \in \mathcal{U}(T)} J^\xi(T, \vec{\pi}, a),$$

and, if it exists, find an admissible strategy ξ^* attaining this value. In this paper we solve (1.3), including giving a full characterization of an optimal control ξ^* and a deterministic numerical method for computing U to arbitrary level of precision. The solution will proceed in two steps: an initial filtering step and a second optimization step. The inference step is studied in Section 2, where we convert the optimal control problem with partial information (1.3) into an equivalent *fully observed* problem in terms of the a posteriori probability process $\vec{\Pi}$. The process $\vec{\Pi}$ summarizes the dynamic updating of controller's beliefs about the Markov chain M given her point process observations. The explicit dynamics of $\vec{\Pi}$ are derived in Proposition 2.2, so that the filtering step is completely solved. The main part of the paper then analyzes the resulting optimal switching problem (2.6) in Sections 3 and 4.

To our knowledge, the finite horizon partially observed switching control problem (which might be viewed as an impulse control problem in terms of ξ) defined in (1.3), has not been studied before. However, it is closely related to optimal stopping problems with partially observable Cox processes that have been extensively looked at starting with the Poisson Disorder problems, see e.g. Peskir and Shiryaev [2000, 2002], Bayraktar and Dayanik [2006], Bayraktar et al. [2006], Bayraktar and Sezer [2006]. In particular, Bayraktar and Sezer [2006] solved the Poisson disorder problem when the change time has phase type prior distribution by showing that it is equivalent to an optimal stopping problem for a hidden Markov process (which has several transient states and one absorbing state) that is indirectly observed through a point process. Later Ludkovski and Sezer [2007] solved a similar optimal stopping problem in which all the states of the hidden Markov chain are recurrent. Both of these works can be viewed as a special case of (1.3), see Remark 3.2. Our model can also be viewed as the continuous-time counterpart of discrete-time sequential M -ary detection in hidden Markov models, a topic extensively studied in sequential analysis, see e.g. Tartakovsky et al. [2006], Aggoun [2003].

Filtering problems with point process observations is a well-studied area; let us mention the work of Arjas et al. [1992], Ceci and Gerardi [1998] and the reference volume Elliott et al. [1995]. In our model we use the previous results obtained in Bayraktar and Sezer [2006], Ludkovski and Sezer [2007] to derive an explicit filter; this allows us then to focus on the separated fully-observed optimal switching problem using the new hyper-state. Let us also mention the recent paper of Chopin and Varini [2007] who study a simulation-based method for filtering in a related model, but where an explicit filter is unavailable and must be numerically approximated.

The techniques that we use to solve the optimal switching/impulse control problem are different from the ones used in the continuous-time optimal control problems mentioned above. The main tool in solving the optimal stopping problems (in the multi-dimensional case, the tools in the one dimensional case are not restricted to the one described here) is the approximating sequence that is constructed by restricting the time horizon to be less than the time of the n -th observation/jump of the observed point process. This sequence converges to the value function uniformly and exponentially fast. However, in the impulse control problem, the corresponding approximating sequence is constructed by restricting the sum of the number of jumps *and* interventions to be less than n . This sequence converges to the value function, however the uniform convergence in both T and $\vec{\pi}$ is not identifiable using the same techniques.

As in Costa and Davis [1989] and Costa and Raymundo [2000] (also see Mazziotto et al. [1988] for general theory of impulse control of partially observed stochastic systems), we first characterize the value function U as the smallest fixed point of two functional operators and obtain the aforementioned approximating sequence. Using one of these characterization results and the path properties of the a posteriori probability process we obtain one of our main contributions: the regularity of the value function U . We show that U is convex in $\vec{\pi}$, Lipschitz in the same variable on

the closure of its domain, and Lipschitz in the T variable uniformly in $\vec{\pi}$. Our regularity analysis leads to the proof of the continuity of U in both T and $\vec{\pi}$ which in turn lets us explicitly describe an optimal strategy.

The other characterization of U as a fixed point of the first jump operator is used to numerically implement the optimal solution and find the value function. In general, very little is known about numerics for continuous-time control of general hidden Markov models, and this implementation is another one of our contributions. We combine the explicit filtering equations together with special properties of piecewise deterministic processes [Davis, 1993] and the structure of general optimal switching problems to give a complete computational scheme. Our method relies only on deterministic optimization sub-problems and lets us avoid having to deal with first order quasi-variational inequalities with integral terms that appear in related stochastic control formulations (see remark 3.3 below). We illustrate our approach with several examples on a finite/infinite horizon and a hidden Markov chain with two or three states.

Our framework has wide-ranging applications in operations research, management science and applied probability. Specific cases are discussed in the next subsection. As these examples demonstrate, our approach leads to sensible policy advice in many scenarios. Most of the relevant applied literature treats discrete-time stationary problems, and our model can be seen as a finite-horizon, continuous-time generalization of these approaches.

The rest of the paper is organized as follows: In Section 1.1 we propose some applications of our modeling framework. In Section 2 we describe an equivalent fully observed problem in terms of the a posteriori probability process $\vec{\Pi}$. We also analyze the dynamics of $\vec{\Pi}$. In Section 3 we show that U satisfies two different dynamic programming equations. The results of Section 3 along with the path description of $\vec{\Pi}$ allows us to study the regularity properties of U and describe an optimal strategy in Section 4. Our model can be extended beyond (1.3), in particular to cover the case of infinite horizon and the case in which the costs are incurred at arrival times. The extensions are described in Section 5. Extensive numerical analysis of several illustrative examples is carried out in Section 6.

1.1. Applications. In this section we discuss case studies of our model and the relevant applied literature.

1.1.1. Cyclical Economic Policy Making. The economic business cycle is a basis of many policy making decisions. For instance, the country's central bank attempts to match its monetary policy, so as to have low interest rates in periods of economic recession and high interest rates when the economy overheats. Similarly, individual firms will time their expenditures to coincide with boom times and will cut back on capital spending in unfavorable economy states. Finally, investors hope to invest in the bull market and stay on the sidelines during the bear market. In all these cases, the precise current economy state is never known. Instead, the agents collect information via

economic events, surveys and news, and act based on their dynamic beliefs about the environment. Typically, such news consist of discrete events (e.g. earnings pre-announcements, geo-political news, economic polls) which cause *instantaneous* jumps in agents' beliefs. Thus, it is natural to model the respective information structure by observations of a modulated compound Poisson process. Accordingly, let M represent the current state of the economy and let the observation X correspond to economic news. Inability to correctly identify M will lead to (opportunity) costs $c_{M_s}(\xi_s)$. Hence, one may take $\mathcal{A} = E$ and $c_a(a) = 0, c_a(b) < 0$. The strategy ξ represents the set of possible actions of the agent. The switching costs of the form $K(\xi_s, \xi_{s-}) > 0$ correspond to the costly influence of the Federal Reserve changing its interest rate policy, or to the transaction costs incurred by the investor who gets in/out of the market. Depending on the particular setting, one may study this problem both in finite- and infinite-horizon setting, and with or without discounting. For instance, a firm planning its capital budgeting expenses might have a fixed horizon of one year, while a central bank has infinite horizon but discounts future costs. A corresponding numerical example is presented in Section 6.2.

1.1.2. Matching Regime-Switching Demand Levels. Many customer-oriented businesses experience stochastically fluctuating demand. Thus, internet servers face heavy/light traffic; manufacturing managers observe cyclical demand levels; customer service centers have varying frequencies of calls. Such systems can be modeled in terms of a compound Poisson request process X modulated by the partially known system state M . Here, X serves the *dual* role of representing the actual demands and conveying information about M . The objective of the agent is to dynamically choose her strategy ξ , so as to track current demand level. For instance, an internet server receives asynchronous requests Y_ℓ , $\ell = 1, 2, \dots$ (corresponding to jumps of X) that take $c(Y_\ell, \xi_t)$ time units to fulfill. The rate of requests and their complexity distribution depend on M . In turn, the server manager can control how much processing power is devoted to the server: more processors cut down individual service times but lead to higher fixed overhead. Such a model effectively corresponds to a controlled $M(\lambda)/G/\infty$ -queue, where the arrival rate λ is M -modulated, and where the distribution of service times depends both on M and the control ξ . A related computational example concerning a customer call center is treated in Section 6.3.

A concrete example that has been recently studied in the literature is the insurance premium problem. Insurance companies handle claims in exchange for policy premiums. A standard model asserts that claims Y_1, Y_2, \dots form a compound (time-inhomogeneous) Poisson process X . Suppose that the rate of claims is driven by some state variable M that measures the current background risk (e.g. climate, health epidemics, etc.), with the latter being unobserved directly. In Aggoun [2003], such a model was studied (in discrete time) from the inference point of view, deriving the optimal filter for the insurance environment M given the claim process. Assume now that the company can control its continuous premium rate $c^2(\xi_t)$, as well as its deductible level $c^1(\xi_t)$.

High deductibles require lowering the premium rate, and are therefore only optimal in high-risk environments. Furthermore, changes to policy provisions (which has a finite expiration date T) are costly and should be undertaken infrequently. The overall objective is thus,

$$\sup_{\xi \in \mathcal{U}(T)} \mathbb{E}^{\vec{\pi}, a} \left[- \sum_{j=1}^{N(T)} e^{-\rho \sigma_j} (Y_j - c^1(\xi_{\sigma_j}))_+ + \int_0^T c^2(\xi_t) dt - \sum_k e^{-\rho \tau_k} \left(\sum_{i \in E} K_i(\xi_{k-1}, \xi_k) \cdot 1_{\{M_{\tau_k}=i\}} \right) \right],$$

where N is the counting process for the number of claims. The resulting cost structure, which is a variant of (1.3), is described in Section 5.2.

1.1.3. Security Monitoring. Classical models of security surveillance (radar, video cameras, communication network monitor) involve an unobserved system state M representing current security (e.g. $E = \{0, 1\}$, where 0 corresponds to a ‘normal’ state and 1 represents a security breach) and a signal X . The signal X records discrete events, namely artifacts in the surveyed space (radar alarms, camera movement, etc.). Benign artifacts are possible, but the intensity λ of X increases when $M_t = 1$. If the signal can be decomposed into further sub-types, then X becomes a marked point process with marks (Y_ℓ). The goal of the monitor is to correctly identify and respond to security breaches, while minimizing false alarms and untreated security violations. Classical formulations [Tartakovsky et al., 2006, Peskir and Shiryaev, 2000] only analyze optimality of the first detection. However, in most practical problems the detection is ongoing and discrete announcement costs require studying the entire (infinite) sequence of detection decisions. Accordingly, our optimal switching framework of (1.3) is more appropriate.

As a simplest case, the monitor can either declare the system to be sound $\xi_t = 1$, or declare a state of alarm $\xi_t = 2$. This produces M -dependent penalty costs at rate $\sum_{j \in E} c_j(\xi_t) 1_{\{M_t=j\}} dt$; also changing the monitor state is costly and leads to costs K . A typical security system is run on an infinite loop and one wishes to minimize total discounted costs, where the discounting parameter ρ models the effective time-horizon of the controller (i.e. the trade-off between the myopically optimal announcement and long-run costs). Such an example is presented in Section 6.1.

1.1.4. Sequential Poisson Disorder Problems. Our model can also serve as a generalization of Poisson disorder problems, [Bayraktar et al., 2006, Peskir and Shiryaev, 2002]. Consider a simple Poisson process X whose intensity λ sequentially alternates between λ_0 and λ_1 . The goal of the observer is to correctly identify the current intensity; doing so produces a running reward at rate $c_0(\xi_t)$ per unit time, otherwise a cost at rate $c_1(\xi_t)$ is assessed, where ξ is the control process. Whenever the observer changes her announcement, a fixed cost K is charged in order to make sure that the agent does not vacillate. Letting $M, M_t \in \{0, 1\}$ denote the intensity state, and $\lambda = \lambda_{M_t}$ this example yet again fits into the framework of (1.3). Obvious generalizations to multiple values of λ and multiple announcement options for the observer can be considered. Again, one may study

the classical infinite-horizon problem, or the harder time-inhomogeneous model on finite-horizon, where the observer must also take into account time-decay costs.

2. PROBLEM STATEMENT

In this section we rigorously define the problem statement and show that it is equivalent to a fully observed impulse control problem using the conditional probability process $\vec{\Pi}$. We then derive explicitly the dynamics of $\vec{\Pi}$. First, however we give a construction of the probability measure \mathbb{P} and the formal description of X .

2.1. Observation Process. Let $(\Omega, \mathcal{H}, \mathbb{P}_0)$ be a probability space hosting two independent elements: (i) a continuous time Markov process M taking values in a finite set E , and with infinitesimal generator $Q = (q_{ij})_{i,j \in E}$, (ii) a compound Poisson process X with intensity λ_1 and jump size distribution ν_1 on \mathbb{R}^d . Let $\mathbb{F} = \{\mathcal{F}_t^X\}$ be the natural filtration of X enlarged by \mathbb{P}_0 -null sets, and consider its initial enlargement $\mathbb{G} = \{\mathcal{G}_t\}_{t \geq 0}$ with $\mathcal{G}_t \triangleq \sigma(\mathcal{F}_t^X, \sigma(\{M_s\}_{s \geq 0}))$ for all $t \geq 0$. The filtration \mathbb{G} summarizes the information flow of a genie that observes the entire path of M at time $t = 0$.

Denote by $\sigma_0, \sigma_1, \dots$ the arrival times of the process X ,

$$\sigma_\ell \triangleq \inf\{t > \sigma_{\ell-1} : X_t \neq X_{t-}\}, \quad \ell \geq 1 \quad \text{with } \sigma_0 \equiv 0.$$

and by Y_1, Y_2, \dots the \mathbb{R}^d -valued marks observed at these arrival times:

$$Y_\ell = X_{\sigma_\ell} - X_{\sigma_{\ell-}}, \quad \ell \geq 1.$$

Then in terms of the counting random measure

$$(2.1) \quad p((0, t], A) \triangleq \sum_{\ell=1}^{\infty} 1_{\{\sigma_\ell \leq t\}} 1_{\{Y_\ell \in A\}},$$

where A is a Borel set in \mathbb{R}^d , we can write the observation process X as

$$X_t = X_0 + \int_{(0,t] \times \mathbb{R}^d} y p(ds, dy).$$

Let us introduce the positive constants $\lambda_2, \dots, \lambda_m$ and the distributions ν_2, \dots, ν_m . We also define the total measure $\nu \triangleq \nu_1 + \dots + \nu_m$, and let $f_i(\cdot)$ be the density of ν_i with respect to ν . Define

$$R(t, y) \triangleq \frac{1}{\lambda_1 f_1(y)} \sum_{i \in E} 1_{\{M_t = i\}} \lambda_i f_i(y), \quad t \geq 0, y \in \mathbb{R}^d.$$

and denote the $(\mathbb{P}_0, \mathbb{G})$ - (or $(\mathbb{P}_0, \mathbb{F})$)-compensator of p by

$$(2.2) \quad p_0((0, t] \times A) = \lambda_1 t \int_A f_1(y) \nu(dy), \quad t \geq 0, A \in \mathcal{B}(\mathbb{R}^d).$$

We will use $R(t, y)$ and p_0 to change the underlying probability measure to a new probability measure \mathbb{P} on (Ω, \mathcal{H}) defined by

$$\frac{d\mathbb{P}}{d\mathbb{P}_0} \Big|_{\mathcal{G}_t} = Z_t,$$

where the stochastic exponential Z given by

$$Z_t \triangleq \exp \left\{ \int_{(0,t] \times \mathbb{R}^d} \log(R(s, y)) p(ds, dy) - \int_{(0,t] \times \mathbb{R}^d} [R(s, y) - 1] p_0(ds, dy) \right\},$$

is a $(\mathbb{P}_0, \mathbb{G})$ -martingale. Note that \mathbb{P} and \mathbb{P}_0 coincide on \mathcal{G}_0 since $Z_0 = 1$, therefore law of the Markov chain M is the same under both probability measures. Moreover, the (\mathbb{P}, \mathbb{G}) -compensator of p becomes

$$(2.3) \quad p_1((0, t], A) = \sum_{i \in E} \int_{(0,t]} 1_{\{M_s=i\}} \lambda_i \int_A f_i(y) \nu(dy) ds.$$

see e.g. [Jacod and Shiryaev \[1987\]](#). The last statement is equivalent to saying that under this new probability, X has the form

$$(2.4) \quad X_t \triangleq X_0 + \int_0^t \sum_{i \in E} 1_{\{M_s=i\}} dX_s^{(i)}, \quad t \geq 0,$$

in which $X^{(1)}, \dots, X^{(m)}$ are independent compound Poisson processes with intensities and jump size distributions $(\lambda_1, \nu_1), \dots, (\lambda_m, \nu_m)$, respectively. Such a process X is called a Markov-modulated Poisson process [\[Karlin and Taylor, 1981\]](#). By construction, the observation process X has independent increments conditioned on $M = \{M_t\}_{t \geq 0}$. Thus, conditioned on $\{M_{\sigma_\ell} = i\}$, the distribution of Y_ℓ is $\nu_i(\cdot)$ on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$.

2.2. Equivalent Fully Observed Problem. Let $D \triangleq \{\vec{\pi} \in [0, 1]^m : \pi_1 + \dots + \pi_m = 1\}$ be the space of prior distributions of the Markov process M . Also, let $\mathcal{S}(s) = \{\tau : \mathbb{F}\text{-stopping time, } \tau \leq s, \mathbb{P}\text{-a.s}\}$ denote the set of all \mathbb{F} -stopping times smaller than or equal to s .

We define the D -valued *conditional probability process* $\vec{\Pi}(t) \triangleq (\Pi_1(t), \dots, \Pi_m(t))$ such that

$$(2.5) \quad \Pi_i(t) = \mathbb{P}\{M_t = i | \mathcal{F}_t^X\}, \quad \text{for } i \in E, \text{ and } t \geq 0.$$

Each component of $\vec{\Pi}$ gives the conditional probability that the current state of M is $\{i\}$ given the information generated by X until the current time t . Using the process $\vec{\Pi}$ we now convert (1.3) into a standard optimal stopping problem.

Proposition 2.1. *The performance of a given strategy $\xi \in \mathcal{U}(T)$ can be written as*

$$(2.6) \quad J^\xi(T, \vec{\pi}, a) = \mathbb{E}^{\vec{\pi}, a} \left[\int_0^T e^{-\rho t} C(\vec{\Pi}(t), \xi_t) dt - \sum_k e^{-\rho \tau_k} K(\xi_{\tau_k-}, \xi_{\tau_k}, \vec{\Pi}(\tau_k)) \right],$$

in terms of the functions

$$(2.7) \quad C(\vec{\pi}, a) \triangleq \sum_{i \in E} c_i(a) \pi_i, \quad \text{and} \quad K(a, b, \vec{\pi}) \triangleq \sum_{i \in E} K_i(a, b) \pi_i.$$

Proposition 2.1 above states that solving the problem in (1.3) is equivalent to solving an impulse control problem with state variables $\vec{\Pi}$ and ξ . As a result, the filtering and optimization steps are completely separated. In our context with optimal switching control, the proof of this separation principle is immediate (see e.g. [Shiryaev \[1978\]](#), pp. 166-167]). In more general problems with continuous controls, the result is more delicate, see [Ceci and Gerardi \[1998\]](#).

We proceed to discuss the technical assumptions on C and K . Note that by construction $C(\cdot, a)$ and $K(a, b, \cdot)$ are linear. Moreover, C is bounded since E is finite, so there is a constant denoted $c = \max_{i \in E} |c_i|$ that uniformly bounds possible rates of profit, $|C(\vec{\pi}, a)| \leq c$. For the switching costs K we assume that they satisfy the triangle inequality

$$K_i(a, b) + K_i(b, c) \geq K_i(a, c), \quad \text{and} \quad K_i(a, b) > k_0 > 0 \quad \text{for } i \in E; a, b, c \in \mathcal{A}.$$

By the above assumptions on the switching costs and because possible rewards are uniformly bounded, with probability one the controller only makes finitely many switches and she does not make two switches at once. Without loss of generality we will also assume that every element in $\xi \in \mathcal{U}(T)$ satisfies

$$(2.8) \quad \mathbb{E}^{\vec{\pi}, a} \left[\sum_k e^{-\rho \tau_k} K(\xi_{\tau_k-}, \xi_{\tau_k}, \vec{\Pi}(\tau_k)) \right] < \infty.$$

Otherwise, the cost associated with a strategy ξ would be $-\infty$ since

$$\mathbb{E}^{\vec{\pi}, a} \left[\int_0^T e^{-\rho t} |C(\vec{\Pi}(t), \xi_t)| dt \right] \leq c T,$$

and taking no action would be better than applying ξ .

In the sequel we will also make use of the following auxiliary problems. First, let U_0 be the value of no-action, i.e.,

$$(2.9) \quad U_0(T, \vec{\pi}, a) = \mathbb{E}^{\vec{\pi}, a} \left[\int_0^T e^{-\rho t} C(\vec{\Pi}_t, a) dt \right].$$

Also in reference to (1.3), we will consider the restricted problems

$$(2.10) \quad U_n(T, \vec{\pi}, a) \triangleq \sup_{\xi \in \mathcal{U}_n(T)} J^\xi(T, \vec{\pi}, a), \quad n \geq 1,$$

in which $\mathcal{U}_n(T)$ is a subset of $\mathcal{U}(T)$ which contains strategies with at most $n \geq 1$ interventions up to time T .

2.3. Sample paths of $\vec{\Pi}$. In this section we describe the filtering procedure of the controller, i.e. the evolution of the conditional probability process $\vec{\Pi}$. Proposition 2.2 explicitly shows that the processes $\vec{\Pi}$ and $(\vec{\Pi}, \xi)$ are piecewise deterministic processes and hence have the strong Markov property, Davis [1993]. This description of paths of the conditional probability process is also discussed in Proposition 2.1 in Ludkovski and Sezer [2007] and Proposition 2.1 of Bayraktar and Sezer [2006]. We summarize the needed results below.

Let

$$(2.11) \quad I(t) \triangleq \int_0^t \sum_{i=1}^m \lambda_i 1_{\{M_s=i\}} ds,$$

so that the probability of no events for the next u time units is $\mathbb{P}^{\vec{\pi}}\{\sigma_1 > u\} = \mathbb{E}^{\vec{\pi}}[e^{-I(u)}]$. Then for $\sigma_\ell \leq t \leq t+u < \sigma_{\ell+1}$, we have

$$(2.12) \quad \Pi_i(t+u) = \left. \frac{\mathbb{P}^{\vec{\pi}}\{\sigma_1 > u, M_u = i\}}{\mathbb{P}^{\vec{\pi}}\{\sigma_1 > u\}} \right|_{\vec{\pi}=\vec{\Pi}(t)}.$$

On the other hand, upon an arrival of size Y_ℓ , the conditional probability $\vec{\Pi}$ experiences a jump

$$(2.13) \quad \Pi_i(\sigma_{\ell+1}) = \frac{\lambda_i f_i(Y_{\ell+1}) \Pi_i(\sigma_{\ell+1}-)}{\sum_{j \in E} \lambda_j f_j(Y_{\ell+1}) \Pi_j(\sigma_{\ell+1}-)}, \quad \text{for } \ell \in \mathbb{N}.$$

To simplify (2.12), define $\vec{x}(t, \vec{\pi}) \equiv (x_1(t, \vec{\pi}), \dots, x_m(t, \vec{\pi}))$ via

$$(2.14) \quad x_i(t, \vec{\pi}) \triangleq \frac{\mathbb{P}^{\vec{\pi}}\{\sigma_1 > t, M_t = i\}}{\mathbb{P}^{\vec{\pi}}\{\sigma_1 > t\}} = \frac{\mathbb{E}^{\vec{\pi}}[1_{\{M_t=i\}} \cdot e^{-I(t)}]}{\mathbb{E}^{\vec{\pi}}[e^{-I(t)}]}, \quad \text{for } i \in E.$$

It can be checked easily that the paths $t \mapsto \vec{x}(t, \vec{\pi})$ have the semigroup property $\vec{x}(t+u, \vec{\pi}) = \vec{x}(u, \vec{x}(t, \vec{\pi}))$. In fact, \vec{x} can be described as a solution of coupled first-order ordinary differential equations. To observe this fact first recall [Darroch and Morris, 1968, Neuts, 1989, Karlin and Taylor, 1981] that the vector

$$(2.15) \quad \vec{m}(t, \vec{\pi}) \equiv (m_1(t, \vec{\pi}), \dots, m_m(t, \vec{\pi})) \triangleq \left(\mathbb{E}^{\vec{\pi}, a}[1_{\{M_t=1\}} \cdot e^{-I(t)}], \dots, \mathbb{E}^{\vec{\pi}, a}[1_{\{M_t=m\}} \cdot e^{-I(t)}] \right)$$

has the form

$$\vec{m}(t, \vec{\pi}) = \vec{\pi} \cdot e^{t(Q-\Lambda)},$$

where Λ is the $m \times m$ diagonal matrix with $\Lambda_{i,i} = \lambda_i$. Thus, the components of $\vec{m}(t, \vec{\pi})$ solve $dm_i(t, \vec{\pi})/dt = -\lambda_i m_i(t, \vec{\pi}) + \sum_{j \in E} m_j(t, \vec{\pi}) \cdot q_{j,i}$ and together with the chain rule and (2.14) we obtain

$$(2.16) \quad \frac{dx_i(t, \vec{\pi})}{dt} = \left(\sum_{j \in E} q_{j,i} x_j(t, \vec{\pi}) - \lambda_i x_i(t, \vec{\pi}) + x_i(t, \vec{\pi}) \sum_{j \in E} \lambda_j x_j(t, \vec{\pi}) \right).$$

For the sequel we note again that $\mathbb{P}^{\vec{\pi}}\{\sigma_1 \in du, M_u = i\} = \mathbb{E}^{\vec{\pi}, a}[\lambda_i 1_{\{M_u=i\}} e^{-I(u)}] du = \lambda_i m_i(u, \vec{\pi}) du$.

The preceding equations (2.12) and (2.13) imply that

Proposition 2.2. *The process $\vec{\Pi}$ is a piecewise-deterministic, (\mathbb{P}, \mathbb{F}) -Markov process. The paths have the characterization*

$$(2.17) \quad \left\{ \begin{array}{l} \vec{\Pi}(t) = \vec{x}\left(t - \sigma_\ell, \vec{\Pi}(\sigma_\ell)\right), \quad \sigma_\ell \leq t < \sigma_{\ell+1}, \quad \ell \in \mathbb{N} \\ \vec{\Pi}(\sigma_\ell) = \left(\frac{\lambda_1 f_1(Y_\ell) \Pi_1(\sigma_\ell-)}{\sum_{j \in E} \lambda_j f_j(Y_\ell) \Pi_j(\sigma_\ell-)}, \dots, \frac{\lambda_m f_m(Y_\ell) \Pi_m(\sigma_\ell-)}{\sum_{j \in E} \lambda_j f_j(Y_\ell) \Pi_j(\sigma_\ell-)} \right) \end{array} \right\}.$$

Alternatively, we can describe $\vec{\Pi}$ in terms of the random measure p ,

$$d\Pi_i(t) = \mu_i(\vec{\Pi}(t-)) dt + \int_{\mathbb{R}^d} J_i(\vec{\Pi}(t-), y) p(dt, dy),$$

for all $i \in E$, where

$$(2.18) \quad \mu_i(\vec{\pi}) = \sum_{j \in E} q_{j,i} \pi_j + \lambda \pi_i \left(\sum_{j \in E} \lambda_j \pi_j - \lambda_i \right), \quad \text{and} \quad J_i(\vec{\pi}, y) = \pi_i \cdot \left(\frac{\lambda_i f_i(y)}{\sum_{j \in E} \lambda_j f_j(y) \pi_j} - 1 \right).$$

Here, one should also note that the (\mathbb{P}, \mathbb{F}) -compensator of the random measure p is

$$\tilde{p}((0, t] \times A) = \sum_{j \in E} \int_0^t \int_A \lambda_j f_j(y) \Pi_j(s) dy ds, \quad t \geq 0, A \text{ Borel}.$$

In more general models with point process observations, an explicit filter for $\vec{\Pi}$ would not be available and one would have to resort to simulation-based approaches, see e.g. [Chopin and Varini \[2007\]](#). The subsequent optimization step would then appear to be intractable, though an integrated Markov chain Monte Carlo paradigm for filtering and optimization was proposed in [Muller et al. \[2004\]](#).

3. TWO DYNAMIC PROGRAMMING EQUATIONS FOR THE VALUE FUNCTION

In this section we establish two dynamic programming equations for the value function U . The first key equation (3.13) reduces the solution of the problem (1.3) to studying a system of coupled optimal stopping problems. The second dynamic programming principle of Proposition 3.4 shows that the value function is also the fixed point of a first jump operator. The latter representation will be useful in the numerical computations.

3.1. Coupled Optimal Stopping Operator. In this section we show that U solves a coupled optimal stopping problem. Combined with regularity results in Section 4, this leads to a direct characterization of an optimal strategy. The analysis of this section parallels the general framework of impulse control of piecewise deterministic processes (pdp) developed by [Costa and Davis \[1989\]](#), [Lenhart and Liao \[1988\]](#). It is also related to optimal stopping of pdp's studied in [Gugerli \[1986\]](#), [Costa and Davis \[1988\]](#).

Let us introduce a functional operator \mathcal{M} whose action on a test function w is

$$(3.1) \quad \mathcal{M}w(T, \vec{\pi}, a) \triangleq \max_{b \in \mathcal{A}, b \neq a} \left\{ w(T, \vec{\pi}, b) - K(a, b, \vec{\pi}) \right\}.$$

The operator \mathcal{M} is called the *intervention* operator and denotes the maximum value that can be achieved if an immediate best change is carried to the current policy. Assuming some ordering on the finite policy set \mathcal{A} , let us denote the smallest policy choice achieving the maximum in (3.1) as

$$(3.2) \quad d_{\mathcal{M}w}(T, \vec{\pi}, a) \triangleq \min_{b \in \mathcal{A}} \left\{ w(T, \vec{\pi}, b) - K(a, b, \vec{\pi}) = \mathcal{M}w(T, \vec{\pi}, a) \right\}.$$

The main object of study in this section is another functional operator \mathcal{G} whose action is described by the following optimal stopping problem:

$$(3.3) \quad \mathcal{G}V(T, \vec{\pi}, a) = \sup_{\tau \in \mathcal{S}(T)} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^\tau e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \tau} \mathcal{M}V(T - \tau, \vec{\Pi}_\tau, a) \right],$$

for $T \in \mathbb{R}_+$, $\vec{\pi} \in D$, and $a \in \mathcal{A}$. We set $V_0 \triangleq U_0$ from (2.9) and iterating \mathcal{G} obtain the following sequence of functions:

$$(3.4) \quad V_{n+1} \triangleq \mathcal{G}V_n, \quad n \geq 0.$$

Lemma 3.1. $(V_n)_{n \in \mathbb{N}}$ is an increasing sequence of functions.

In Section 4 we will further show that (V_n) are convex and continuous.

Proof. The statement follows since

$$\begin{aligned} V_1(T, \vec{\pi}, a) &= \mathcal{G}V_0(T, \vec{\pi}, a) = \sup_{\tau \in \mathcal{S}(T)} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^\tau e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \tau} \mathcal{M}V_0(T - \tau, \vec{\Pi}_\tau, a) \right] \\ &\geq \mathbb{E}^{\vec{\pi}, a} \left[\int_0^T e^{-\rho s} C(\vec{\Pi}_s, a) ds \right] = U_0(T, \vec{\pi}, a) = V_0(T, \vec{\pi}, a), \end{aligned}$$

and since \mathcal{G} is a monotone/positive operator, i.e. for any two functions $f_1 \leq f_2$ we have $\mathcal{G}f_1 \leq \mathcal{G}f_2$, and \square

The following proposition shows that the value functions $(U_n)_{n \in \mathbb{N}}$ of (2.10), which correspond to the restricted control problems over $\mathcal{U}_n(T)$, can be alternatively obtained via the sequence of iterated optimal stopping problems in (3.4).

Proposition 3.1. $U_n = V_n$ for $n \in \mathbb{N}$.

Proof. By definition we have that $U_0 = V_0$. Let us assume that $U_n = V_n$ and show that $U_{n+1} = V_{n+1}$. We will carry out the proof in two steps.

Step 1. First we will show that $U_{n+1} \leq V_{n+1}$. Let $\xi \in \mathcal{U}_{n+1}(T)$,

$$\xi_t = \sum_{k=0}^{n+1} \xi_k \cdot 1_{[\tau_k, \tau_{k+1})}(t), \quad t \in [0, T],$$

with $\tau_0 = 0$ and $\tau_{n+1} = T$, be ε -optimal for the problem in (2.10), i.e.,

$$(3.5) \quad U_{n+1}(T, \vec{\pi}, a) - \varepsilon \leq J^\xi(T, \vec{\pi}, a).$$

Let $\tilde{\xi} \in \mathcal{U}_n(T)$ be defined as

$$\tilde{\xi}_t = \sum_{k=0}^n \tilde{\xi}_k \cdot 1_{[\tilde{\tau}_k, \tilde{\tau}_{k+1})}(t), \quad t \in [0, T],$$

in which $\tilde{\tau}_0 = 0$, $\tilde{\xi}_0 = a$, and $\tilde{\tau}_n = \tau_{n+1}$, $\tilde{\xi}_n = \xi_{n+1}$, for $n \in \mathbb{N}_+$. Using the strong Markov property of $(\vec{\Pi}, \xi)$, we can write J^ξ as

$$(3.6) \quad \begin{aligned} J^\xi(T, \vec{\pi}, a) &= \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\tau_1} e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \tau_1} \left(J^{\tilde{\xi}}(T - \tau_1, \vec{\Pi}_{\tau_1}, \xi_1) - K(a, \xi_1, \vec{\Pi}_{\tau_1}) \right) \right] \\ &\leq \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\tau_1} e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \tau_1} \left(V_n(T - \tau_1, \vec{\Pi}_{\tau_1}, \xi_1) - K(a, \xi_1, \vec{\Pi}_{\tau_1}) \right) \right] \\ &\leq \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\tau_1} e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \tau_1} \mathcal{M}V_n(T - \tau_1, \vec{\Pi}_{\tau_1}, \xi_1) \right] \\ &\leq \mathcal{G}V_n(T, \vec{\pi}, a) = V_{n+1}(T, \vec{\pi}, a). \end{aligned}$$

Here, the first inequality follows from induction hypothesis, the second inequality follows from the definition of \mathcal{M} , and the last inequality from the definition of \mathcal{G} . As a result of (3.5) and (3.6) we have that $U_{n+1} \leq V_{n+1}$ since $\varepsilon > 0$ is arbitrary.

Step 2. To show the opposite inequality $U_{n+1} \geq V_{n+1}$, we will construct a special $\bar{\xi} \in \mathcal{U}_{n+1}(T)$. To this end let us introduce

$$(3.7) \quad \begin{cases} \bar{\tau}_1 = \inf\{t \geq 0 : \mathcal{M}V_n(T - t, \vec{\Pi}_t, a) \geq V_{n+1}(T - t, \vec{\Pi}_t, a) - \varepsilon\}, \\ \bar{\xi}_1 = d_{\mathcal{M}V_n}(T - \bar{\tau}_1, \vec{\Pi}_{\bar{\tau}_1}, a). \end{cases}$$

Let $\hat{\xi}_t = \sum_{k=0}^n \hat{\xi}_k \cdot 1_{[\hat{\tau}_k, \hat{\tau}_{k+1})}(t)$, $\hat{\xi} \in \mathcal{U}_n(T)$ be ε -optimal for the problem in which n interventions are allowed, i.e. (2.10). Using $\hat{\xi}$ we now complete the description of the control $\bar{\xi} \in \mathcal{U}_{n+1}(T)$ by assigning,

$$(3.8) \quad \bar{\tau}_{n+1} = \hat{\tau}_n \circ \theta_{\tau_1}, \quad \bar{\xi}_{n+1} = \hat{\xi}_n \circ \theta_{\bar{\tau}_1}, \quad n \in \mathbb{N}_+,$$

in which θ is the classical *shift operator* used in the theory of Markov processes.

Note that $\bar{\tau}_1$ is an ε -optimal stopping time for the stopping problem in the definition of $\mathcal{G}V_n$. This follows from the classical optimal stopping theory since the process $\vec{\Pi}$ has the strong Markov

property. Therefore,

$$(3.9) \quad \begin{aligned} V_{n+1}(T, \vec{\pi}, a) - \varepsilon &\leq \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\bar{\tau}_1} e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \bar{\tau}_1} \mathcal{M}V_n(T - \bar{\tau}_1, \vec{\Pi}_{\bar{\tau}_1}, a) \right] \\ &\leq \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\bar{\tau}_1} e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \bar{\tau}_1} \left(U_n(T - \bar{\tau}_1, \vec{\Pi}_{\bar{\tau}_1}, \bar{\xi}_1) - K(a, \bar{\xi}_1, \vec{\Pi}_{\bar{\tau}_1}) \right) \right], \end{aligned}$$

in which the second inequality follows from the definition of $\bar{\xi}_1$ and the induction hypothesis. It follows from (3.9) and the strong Markov property of $(\vec{\Pi}, \xi)$ that

$$(3.10) \quad \begin{aligned} V_{n+1}(T, \vec{\pi}, a) - 2\varepsilon &\leq \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\bar{\tau}_1} e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \bar{\tau}_1} \left(U_n(T - \bar{\tau}_1, \vec{\Pi}_{\bar{\tau}_1}, \bar{\xi}_1) - \varepsilon - K(a, \bar{\xi}_1, \vec{\Pi}_{\bar{\tau}_1}) \right) \right] \\ &\leq \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\bar{\tau}_1} e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \bar{\tau}_1} \left(J^{\hat{\xi}}(T - \bar{\tau}_1, \vec{\Pi}_{\bar{\tau}_1}, \xi_1) - K(a, \xi_1, \vec{\Pi}_{\bar{\tau}_1}) \right) \right] \\ &= J^{\bar{\xi}}(T, \vec{\pi}, a) \leq U_{n+1}(T, \vec{\pi}, a). \end{aligned}$$

This completes the proof of the second step since $\varepsilon > 0$ is arbitrary. \square

Proposition 3.2. $\lim_{n \uparrow \infty} V_n(T, \vec{\pi}, a) = U(T, \vec{\pi}, a)$, for any $T \in \mathbb{R}_+$, $\vec{\pi} \in D$, $a \in \mathcal{A}$.

Proof. Fix $(T, \vec{\pi}, a)$. The monotone limit $V(T, \vec{\pi}, a) = \lim_{n \rightarrow \infty} V_n(T, \vec{\pi}, a)$ exists as a result of Lemma 3.1. Since $\mathcal{U}_n(T) \subset \mathcal{U}(T)$, it follows that $V_n(T, \vec{\pi}, a) = U_n(T, \vec{\pi}, a) \leq U(T, \vec{\pi}, a)$. Therefore $V(T, \vec{\pi}, a) \leq U(T, \vec{\pi}, a)$. In the remainder of the proof we will show that $V(T, \vec{\pi}, a) \geq U(T, \vec{\pi}, a)$.

Let $\xi \in \mathcal{U}(T)$ be given, and let $\tilde{\xi}_t := \xi_{t \wedge \tau_n}$, $\tilde{\xi} \in \mathcal{U}_n(T)$, correspond to ξ up to its n -th switch. Then

$$(3.11) \quad \begin{aligned} |J^\xi(T, \vec{\pi}, a) - J^{\tilde{\xi}}(T, \vec{\pi}, a)| &\leq \mathbb{E}^{\vec{\pi}, a} \left[\int_{\tau_n}^T e^{-\rho s} |C(\vec{\Pi}_s, \xi_s) - C(\vec{\Pi}_s, \tilde{\xi}_{\tau_n})| ds + \sum_{k \geq n+1} e^{-\rho \tau_k} K(\xi_{\tau_{k-1}}, \xi_{\tau_k}, \vec{\Pi}_{\tau_k}) \right]. \end{aligned}$$

Now, the right-hand-side of (3.11) converges to 0 as $n \rightarrow \infty$: on the one hand observe that by monotone convergence theorem and (2.8)

$$\lim_{n \rightarrow \infty} \mathbb{E}^{\vec{\pi}, a} \left[\sum_{k \geq n+1} e^{-\rho \tau_k} K(\xi_{\tau_{k-1}}, \xi_{\tau_k}, \vec{\Pi}_{\tau_k}) \right] = 0.$$

On the other hand, since there are only finitely many switches almost surely for any given path,

$$\lim_{n \rightarrow \infty} \int_0^T 1_{\{s > \tau_n\}} e^{-\rho s} |C(\vec{\Pi}_s, \xi_s) - C(\vec{\Pi}_s, \tilde{\xi}_{\tau_n})| ds = 0,$$

and $\int_{\tau_n}^T e^{-\rho s} |C(\vec{\Pi}_s, \xi_s) - C(\vec{\Pi}_s, \xi_{\tau_n})| ds \leq 2cT$. Therefore, the dominated convergence theorem implies that

$$\lim_{n \rightarrow \infty} \mathbb{E}^{\vec{\pi}, a} \left[\int_{\tau_n}^T e^{-\rho s} |C(\vec{\Pi}_s, \xi_s) - C(\vec{\Pi}_s, \tilde{\xi}_{\tau_n})| ds \right] = 0.$$

As a result, for any $\varepsilon > 0$ and n large enough, we find

$$|J^\xi(T, \vec{\pi}, a) - J^{\tilde{\xi}}(T, \vec{\pi}, a)| \leq \varepsilon.$$

Now, since $\tilde{\xi} \in \mathcal{U}_n(T)$ we have $V_n(T, \vec{\pi}, a) = U_n(T, \vec{\pi}, a) \geq J^{\tilde{\xi}}(T, \vec{\pi}, a) \geq J^\xi(T, \vec{\pi}, a) - \varepsilon$ for sufficiently large n , and it follows that

$$(3.12) \quad V(T, \vec{\pi}, a) = \lim_{n \rightarrow \infty} V_n(T, \vec{\pi}, a) \geq J^\xi(T, \vec{\pi}, a) - \varepsilon.$$

Since ξ and ε are arbitrary, we have the desired result. \square

Proposition 3.3. *The value function U is the smallest solution of the dynamic programming equation $\mathcal{G}U = U$, such that $U \geq U_0$. Thus,*

$$(3.13) \quad U(T, \vec{\pi}, a) = \sup_{\tau \in \mathcal{S}(T)} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^\tau e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \tau} \mathcal{M}U(T - \tau, \vec{\Pi}_\tau, a) \right].$$

Proof. **Step 1.** First we will show that U is a fixed point of \mathcal{G} . Since $V_n \leq U$, monotonicity of \mathcal{G} implies that

$$V_{n+1}(T, \vec{\pi}, a) \leq \sup_{\tau \in \mathcal{S}(T)} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^\tau e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \tau} \mathcal{M}U(T - \tau, \vec{\Pi}_\tau, a) \right].$$

Taking the limit of the left-hand-side with respect to n and using Lemma 3.1 and Proposition 3.2 we have

$$U(T, \vec{\pi}, a) \leq \sup_{\tau \in \mathcal{S}(T)} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^\tau e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \tau} \mathcal{M}U(T - \tau, \vec{\Pi}_\tau, a) \right].$$

Let us obtain the reverse inequality. Let $\tilde{\tau} \in \mathcal{S}(T)$ be an ε -optimal stopping time for the problem in the definition of $\mathcal{G}U$, i.e.,

$$(3.14) \quad \begin{aligned} & \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\tilde{\tau}} e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \tilde{\tau}} \mathcal{M}U(T - \tilde{\tau}, \vec{\Pi}_{\tilde{\tau}}, a) \right] \\ & \geq \sup_{\tau \in \mathcal{S}(T)} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^\tau e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \tau} \mathcal{M}U(T - \tau, \vec{\Pi}_\tau, a) \right] - \varepsilon. \end{aligned}$$

Then, as a result of monotone convergence theorem and Proposition 3.2

(3.15)

$$\begin{aligned} U(T, \vec{\pi}, a) &= \lim_{n \rightarrow \infty} V_n(T, \vec{\pi}, a) \geq \lim_{n \rightarrow \infty} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\tilde{\tau}} e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \tilde{\tau}} \mathcal{M}V_{n-1}(T - \tilde{\tau}, \vec{\Pi}_{\tilde{\tau}}, a) \right] \\ &= \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\tilde{\tau}} e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \tilde{\tau}} \mathcal{M}U(T - \tilde{\tau}, \vec{\Pi}_{\tilde{\tau}}, a) \right]. \end{aligned}$$

Now, (3.14) and (3.15) together yield the desired result since ε is arbitrary.

Step 2. Let \tilde{U} be another fixed point of \mathcal{G} satisfying $\tilde{U} \geq U_0 = V_0$. Then an induction argument shows that $\tilde{U} \geq U$: assume that $\tilde{U} \geq V_n$. Then $\mathcal{G}\tilde{U} \geq \mathcal{G}V_n = V_{n+1}$, by the monotonicity of \mathcal{G} . Therefore for all n , $\tilde{U} \geq V_n$, which implies that $\tilde{U} \geq \sup_n V_n = U$. \square

To illustrate the nature of (3.13) consider the special case where $\mathcal{A} = \{1, 2\}$ so that only two types of policies are available. In that case the intervention operator \mathcal{M} is trivial, $\mathcal{M}U(t, \vec{\pi}, a) = U(t, \vec{\pi}, 3 - a) - K(a, 3 - a, \vec{\pi})$. For ease of notation we write $U(t, \vec{\pi}, 1) =: V(t, \vec{\pi})$, $U(t, \vec{\pi}, 2) =: W(t, \vec{\pi})$. It follows that (3.13) can be written as two coupled optimal stopping problems:

$$\begin{cases} V(T, \vec{\pi}) = \sup_{\tau \in \mathcal{S}(T)} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^\tau e^{-\rho s} C(\vec{\Pi}_s, 1) ds + e^{-\rho \tau} (W(T - \tau, \vec{\Pi}_\tau) - K(1, 2, \vec{\pi})) \right] \\ W(T, \vec{\pi}) = \sup_{\tau \in \mathcal{S}(T)} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^\tau e^{-\rho s} C(\vec{\Pi}_s, 2) ds + e^{-\rho \tau} (V(T - \tau, \vec{\Pi}_\tau) - K(2, 1, \vec{\pi})) \right]. \end{cases}$$

The next section discusses how to solve such coupled systems.

Remark 3.1. The value function $U(T, \cdot, a)$ is uniformly bounded. Indeed,

$$U(T, \vec{\pi}, a) \geq U_0(T, \vec{\pi}, a) = \mathbb{E}^{\vec{\pi}, a} \left[\int_0^T e^{-\rho s} C(\vec{\Pi}_s, a) ds \right] \geq - \int_0^T e^{-\rho s} c ds,$$

and conversely for any $\xi \in \mathcal{U}(T)$,

$$J^\xi(T, \vec{\pi}, a) \leq \mathbb{E}^{\vec{\pi}, a} \left[\int_0^T e^{-\rho s} C(\vec{\Pi}_s, \xi_s) ds \right] \leq \int_0^T e^{-\rho s} c ds.$$

Since

$$\int_0^T e^{-\rho s} c ds \leq \begin{cases} cT & \text{when } \rho = 0; \\ c/\rho & \text{when } \rho > 0, \end{cases}$$

we see that when $\rho > 0$ those bounds are even uniform in T .

Remark 3.2. One may extend the above analysis to cover the slightly more general case where $K(a, b, \vec{\pi})$ are allowed to be negative, as long as we assume that for *any* chain $a_0, a_1, \dots, a_n, a_i \in \mathcal{A}$ we have

$$K(a_0, a_1, \vec{\pi}) + K(a_1, a_2, \vec{\pi}) + \dots + K(a_n, a_0, \vec{\pi}) > k_0 > 0,$$

uniformly. This condition implies that repeated switching is unprofitable and guarantees that the number of switches along any path is finite with probability one. Then taking $\mathcal{A}' = \{0, \Delta_1, \dots, \Delta_A\}$ and for any $i \in \mathcal{A}$, $K(0, \Delta_i, \vec{\pi}) = -\sum_{j \in E} H(i, j)\pi_j$, $K(\Delta_i, 0, \vec{\pi}) = +\infty$, $C(\vec{\pi}, \Delta_i) = 0$, one may

imbed the optimal stopping problems studied in Bayraktar and Sezer [2006] and Ludkovski and Sezer [2007] in our framework. Namely, it is easy to see that in this case

$$(3.16) \quad U(T, \vec{\pi}, 0) = \sup_{\tau \in \mathcal{S}(T), \xi_1 \in \mathcal{A}} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^\tau e^{-\rho s} C(\vec{\Pi}_s, 0) ds + e^{-\rho \tau} H(\xi_1, M_\tau) \right].$$

In that sense, our model is a direct extension of optimal stopping problems for hidden Markov models with Poissonian observations.

Remark 3.3. Using the dynamic programming principle developed in Proposition 3.3 one expects that the value function U is the unique weak solution of a coupled system of QVIs (quasi-variational inequalities)

$$(3.17) \quad \begin{aligned} -\frac{\partial}{\partial T} U(T, \vec{\pi}, a) + \mathcal{A}U(T, \vec{\pi}, a) - \rho U(T, \vec{\pi}, a) + C(\vec{\pi}, a) &\leq 0 \\ U(T, \vec{\pi}, a) &\geq \mathcal{M}U(T, \vec{\pi}, a) \\ \left(-\frac{\partial}{\partial T} U(T, \vec{\pi}, a) + \mathcal{A}U(T, \vec{\pi}, a) - \rho U(T, \vec{\pi}, a) + C(\vec{\pi}, a) \right) (U((T, \vec{\pi}, a)) - \mathcal{M}U(T, \vec{\pi}, a)) &= 0. \end{aligned}$$

Here \mathcal{A} is the infinitesimal generator of the process $\vec{\Pi}$ given by (2.9). \mathcal{A} is a first order integro-differential operator. Note that the differential operators do not differentiate with respect to a , therefore for each a we obtain a different QVI. These QVIs are coupled by the action of the intervention operator \mathcal{M} .

One could attempt to numerically solve the above system of QVIs. However, the theoretical basis for the QVI formulation requires justification, in particular in terms of the regularity of the value function U . Typically one must pass to the realm of viscosity solutions to make progress; in contrast in the next section we will develop another a more direct characterization of the value function (see Proposition 3.4). In Section 4 we will use this characterization to develop the regularity properties of U , which helps us describe an optimal control. The more direct characterization of the value function in Proposition 3.4 also provides us a numerical method for numerically solving for the value function.

3.2. First Jump Operator. The following Proposition 3.4 shows that the value function U satisfies a second dynamic programming principle, namely U is the fixed point of the first jump operator \hat{L} . This representation will be used in our numerical computations in Section 6. Let us introduce a functional operator L whose action on test functions V and H is given by

$$(3.18) \quad \begin{aligned} L(V, H)(T, \vec{\pi}, a) &= \sup_{t \in [0, T]} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{t \wedge \sigma_1} e^{-\rho s} C(\vec{\Pi}_s, a) ds \right. \\ &\quad \left. + 1_{\{t < \sigma_1\}} e^{-\rho t} H(T - t, \vec{\Pi}_t, a) + e^{-\rho \sigma_1} 1_{\{t \geq \sigma_1\}} V(T - \sigma_1, \vec{\Pi}_{\sigma_1}, a) \right]. \end{aligned}$$

Observe that L is clearly monotone in both of its function arguments. Moreover, we have

$$(3.19) \quad \begin{aligned} L(V, H)(T, \vec{\pi}, a) &= \sup_{\tau \in \mathcal{S}(T)} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\tau \wedge \sigma_1} e^{-\rho s} C(\vec{\Pi}_s, a) ds + 1_{\{\tau < \sigma_1\}} e^{-\rho t} H(T - \tau, \vec{\Pi}_\tau, a) \right. \\ &\quad \left. + e^{-\rho \sigma_1} 1_{\{\tau \geq \sigma_1\}} V(T - \sigma_1, \vec{\Pi}_{\sigma_1}, a) \right], \end{aligned}$$

which follows as a result of the characterization of the stopping times of piecewise deterministic Markov processes (Theorem T.33 Bremaud [1981], and Theorem A2.3 Davis [1993]) which state that for any $\tau \in \mathcal{S}(T)$, $\tau \wedge \sigma_1 = t \wedge \sigma_1$ for some constant t .

Let us introduce another monotone functional operator by

$$\hat{L}V \triangleq L(V, \mathcal{M}V).$$

Proposition 3.4. *U is the smallest fixed point of \hat{L} that is larger than U_0 . Moreover, the following sequence which is constructed by iterating \hat{L} ,*

$$(3.20) \quad W_0 \triangleq U_0, \quad W_{n+1} \triangleq \hat{L}W_n, \quad n \in \mathbb{N},$$

satisfies $W_n \nearrow U$ (pointwise).

Proof. **Step 1.** Recall that \hat{L} is a monotone operator and that

$$\begin{aligned} W_1(T, \vec{\pi}, a) &= L(U_0, \mathcal{M}U_0)(T, \vec{\pi}, a) \geq \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{T \wedge \sigma_1} e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \sigma_1} 1_{\{T \geq \sigma_1\}} U_0(T - \sigma_1, \vec{\Pi}_{\sigma_1}, a) \right] \\ &= U_0(T, \vec{\pi}, a) = W_0(T, \vec{\pi}, a). \end{aligned}$$

Therefore $(W_n)_{n \in \mathbb{N}}$ is an increasing sequence of functions. Denote the pointwise limit of this sequence by $W = \sup_n W_n$. This limit is a fixed point of \hat{L} :

$$(3.21) \quad \begin{aligned} W(T, \vec{\pi}, a) &= \sup_{n \in \mathbb{N}} W_n(T, \vec{\pi}, a) \\ &= \sup_{n \in \mathbb{N}} \sup_{t \in [0, T]} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{t \wedge \sigma_1} e^{-\rho s} C(\vec{\Pi}_s, a) ds + 1_{\{t < \sigma_1\}} e^{-\rho t} \mathcal{M}W_{n-1}(T - t, \vec{\Pi}_t, a) \right. \\ &\quad \left. + e^{-\rho \sigma_1} 1_{\{t \geq \sigma_1\}} W_{n-1}(T - \sigma_1, \vec{\Pi}_{\sigma_1}, a) \right] \\ &= \sup_{t \in [0, T]} \sup_{n \in \mathbb{N}} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{t \wedge \sigma_1} e^{-\rho s} C(\vec{\Pi}_s, a) ds + 1_{\{t < \sigma_1\}} e^{-\rho t} \mathcal{M}W_{n-1}(T - t, \vec{\Pi}_t, a) \right. \\ &\quad \left. + e^{-\rho \sigma_1} 1_{\{t \geq \sigma_1\}} W_{n-1}(T - \sigma_1, \vec{\Pi}_{\sigma_1}, a) \right] = \hat{L}W(T, \vec{\pi}, a), \end{aligned}$$

where the last line follows from the monotone convergence theorem. In fact it is the smallest of the fixed points of \hat{L} that is greater than $U_0 = W_0$, which is a result of the following induction argument: suppose that $\tilde{W} \geq U_0$ is another such fixed point. Then $\tilde{W} = \hat{L}\tilde{W} \geq \hat{L}U_0 = W_1$. On

the other hand, if $\tilde{W} \geq W_n$, then $\tilde{W} = \hat{L}\tilde{W} \geq \hat{L}W_n = W_{n+1}$. Now taking the supremum of both sides we have that $\tilde{W} \geq W$.

Step 2. We will now show that W is a fixed point of \mathcal{G} , hence $W \geq U$ as a result of Proposition 3.3. First, we will show that $W \geq \mathcal{G}W$. Let us construct an increasing sequence of functions by $u_0 = \mathcal{M}W$, $u_{n+1} = L(u_n, \mathcal{M}W)$, $n \in \mathbb{N}$. It can be shown that u_n can be written as

$$(3.22) \quad u_n(T, \vec{\pi}, a) = \sup_{\tau \in \mathcal{S}(T)} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\tau \wedge \sigma_n} e^{-\rho s} C(\vec{\Pi}_s, a) ds + e^{-\rho \tau \wedge \sigma_n} \mathcal{M}W(T - (\tau \wedge \sigma_n), \vec{\Pi}_{\tau \wedge \sigma_n}, a) \right],$$

see e.g. Proposition 5.5 in Bayraktar et al. [2006]. Taking $n \rightarrow \infty$ we find that the monotone limit $u = \lim_{n \uparrow \infty} u_n$ satisfies $u = \mathcal{G}W$. Now, we can show that $W \geq \mathcal{G}W$ using induction. From step 1, we know that $W = L(W, \mathcal{M}W)$, therefore $W \geq \mathcal{M}W = u_0$ (since stopping immediately may not be optimal in (3.19)). On the other hand, if $W \geq u_n$, then since $L(\cdot, \mathcal{M}W)$ is a monotone operator, we have that $W = L(W, \mathcal{M}W) \geq L(u_n, \mathcal{M}W) = u_{n+1}$. This implies that $W \geq u_n$ for all $n \in \mathbb{N}$. Therefore, $W \geq \mathcal{G}W = \sup_n u_n$.

Let us show the reverse inequality: $W \leq \mathcal{G}W$. As a result of the monotone convergence theorem we have that $\mathcal{G}W = \sup_{n \in \mathbb{N}} \mathcal{G}W_n$. Clearly $\mathcal{G}W_n \geq \hat{L}W_n$ since $W_n \geq \mathcal{M}W_{n-1}$, and the set of stopping times that we are taking a sup over is smaller than $\mathcal{S}(T)$. Therefore, $\mathcal{G}W_n \geq W_{n+1}$. Since we can repeat this argument for all n ,

$$\mathcal{G}W = \sup_{n \in \mathbb{N}} \mathcal{G}W_n \geq \sup_{n \in \mathbb{N}} W_{n+1} = W.$$

Step 3. We will now show that $W \leq U$ (which together with the result of step 2, shows that $W = U$). On the one hand, using the strong Markov property of $(\vec{\Pi}, \xi)$, the value function U can be shown to be a fixed point of \hat{L} (see Proposition 5.6 in Bayraktar et al. [2006]): recall that $U = \mathcal{G}U$ (the right-hand-side of which is an optimal stopping problem) and compare with (3.19). On the other hand, from step 1 we know that W is the smallest fixed point of \hat{L} greater than U_0 . But this implies that $U \geq W$. \square

Remark 3.4. As a result of Fubini's theorem and using (2.13) and (2.15) we can write \hat{L} as

$$(3.23) \quad \hat{L}V(T, \vec{\pi}, a) = \sup_{0 \leq t \leq T} \left\{ \left(\sum_{i \in E} m_i(t, \vec{\pi}) \right) \cdot e^{-\rho t} \mathcal{M}V(T - t, \vec{x}(t, \vec{\pi}), a) \right. \\ \left. + \int_0^t e^{-\rho u} \sum_{i \in E} m_i(u, \vec{\pi}) \cdot \left(C(\vec{x}(u, \vec{\pi}), a) + \lambda_i \cdot S_i V(T - u, \vec{x}(u, \vec{\pi}), a) \right) du \right\},$$

in terms of the operator

$$(3.24) \quad S_i w(t, \vec{\pi}, a) \triangleq \int_{\mathbb{R}^d} w \left(t, \left(\frac{\lambda_1 f_1(y) \pi_1}{\sum_{j \in E} \lambda_j f_j(y) \pi_j}, \dots, \frac{\lambda_m f_m(y) \pi_m}{\sum_{j \in E} \lambda_j f_j(y) \pi_j} \right), a \right) f_i(y) \nu(dy), \quad \text{for } i \in E.$$

This implies that one can numerically compute $\hat{L}V$ by performing the *deterministic* optimization on the right-hand-side of (3.23).

4. REGULARITY OF THE VALUE FUNCTION AND AN OPTIMAL STRATEGY

In this section we will analyze the regularity of the value function U , which will lead to the construction of an optimal strategy. This is done by analysis of two auxiliary sequences of functions converging to U . We first begin by studying U_0 .

Lemma 4.1. *The function U_0 defined in (2.9) is convex in $\vec{\pi}$.*

Proof. Let us define a functional operator I through its action on a test function w by $Iw = L(w, 0)$, that is,

$$(4.1) \quad \begin{aligned} Iw(T, \vec{\pi}, a) &= \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\sigma_1 \wedge T} e^{-\rho t} C(\vec{\Pi}_t, a) dt + 1_{\{\sigma_1 \leq T\}} e^{-\rho \sigma_1} w(T - \sigma_1, \vec{\Pi}_{\sigma_1}, a) \right] \\ &= \int_0^T e^{-\rho u} \sum_{i \in E} m_i(u, \vec{\pi}) \cdot [C(\vec{x}(u, \vec{\pi}), a) + \lambda_i \cdot S_i w(T - u, \vec{x}(u, \vec{\pi}), a)] du. \end{aligned}$$

As a result of the strong Markov property of $\vec{\Pi}$ we observe that U_0 is a fixed point of I , and if we define

$$(4.2) \quad k_{n+1}(T, \vec{\pi}, a) = Ik_n(T, \vec{\pi}, a), \quad k_0(T, \vec{\pi}, a) = 0, \quad T \in \mathbb{R}_+, \vec{\pi} \in D, a \in \mathcal{A}$$

then $k_n \nearrow U_0$, see Proposition 1 in Costa and Davis [1989]. We will divide the rest of the proof into two parts. In the first part we will show that k_n converges to U_0 uniformly. In the second part we will argue that for all $n \in \mathbb{N}$, k_n is convex. Suppose both of the above claims have been proved and let $\varepsilon > 0$. Then for any $\vec{\pi}_1, \vec{\pi}_2 \in D$

$$(4.3) \quad \begin{aligned} U_0(T, \alpha \vec{\pi}_1 + (1 - \alpha) \vec{\pi}_2, a) &= U_0(T, \alpha \vec{\pi}_1 + (1 - \alpha) \vec{\pi}_2, a) - k_n(T, \alpha \vec{\pi}_1 + (1 - \alpha) \vec{\pi}_2, a) \\ &\quad + k_n(T, \alpha \vec{\pi}_1 + (1 - \alpha) \vec{\pi}_2, a) \\ &\leq \varepsilon + \alpha k_n(T, \vec{\pi}_1, a) + (1 - \alpha) k_n(T, \vec{\pi}_2, a) \\ &\leq 2\varepsilon + \alpha U_0(T, \vec{\pi}_1, a) + (1 - \alpha) U_0(T, \vec{\pi}_2, a), \end{aligned}$$

in which the last two inequalities follow since for $n > N(\varepsilon)$ large enough, $|U_0(T, \vec{\pi}, a) - k_n(T, \vec{\pi}, a)| < \varepsilon$ for all $\vec{\pi} \in D$. Since ε was arbitrary the convexity of $\vec{\pi} \rightarrow U_0(T, \vec{\pi}, a)$ follows.

Step 1. Using strong Markov property we can write k_n as (cf. (3.22))

$$(4.4) \quad k_n(T, \vec{\pi}, a) = \mathbb{E}^{\vec{\pi}} \left[\int_0^{\sigma_n \wedge T} e^{-\rho t} C(\vec{\Pi}_t, a) dt \right].$$

As a result,

$$\begin{aligned}
|U_0(T, \vec{\pi}, a) - k_n(T, \vec{\pi}, a)| &\leq \mathbb{E}^{\vec{\pi}} \left[1_{\{T > \sigma_n\}} \int_{\sigma_n}^T e^{-\rho t} |C(\vec{\Pi}_t, a)| dt \right] \\
(4.5) \quad &\leq \mathbb{E}^{\vec{\pi}} \left[1_{\{T > \sigma_n\}} e^{-\rho \sigma_n} c \int_0^{T-\sigma_n} e^{-\rho t} dt \right] \\
&\leq c T \mathbb{P}^{\vec{\pi}} \{T > \sigma_n\} \\
&\leq c T \mathbb{E}^{\vec{\pi}} [1_{\{T > \sigma_n\}} (T/\sigma_n)] \leq c T^2 \cdot \mathbb{E}^{\vec{\pi}} [1/\sigma_n].
\end{aligned}$$

The conditional probability of the first jump satisfies $\mathbb{P}^{\vec{\pi}} \{\sigma_1 > t|M\} = e^{-I(t)}$. Therefore,

$$\begin{aligned}
\mathbb{E}^{\vec{\pi}} [e^{-u\sigma_1} | M] &= \mathbb{E}^{\vec{\pi}} \left[\int_{\sigma_1}^{\infty} ue^{-ut} , dt \mid M \right] = \int_0^{\infty} \mathbb{P}^{\vec{\pi}} \{\sigma_1 \leq t|M\} \cdot ue^{-ut} dt \\
(4.6) \quad &= \int_0^{\infty} [1 - e^{-I(t)}] ue^{-ut} dt \\
&\leq \int_0^{\infty} [1 - e^{-\bar{\lambda}t}] ue^{-ut} dt = \frac{\bar{\lambda}}{\bar{\lambda} + u},
\end{aligned}$$

where $\bar{\lambda} = \max_{i \in E} \lambda_i$, see (2.11). Since the observed process X has independent increments given M , it readily follows that $\mathbb{E}^{\vec{\pi}} [e^{-u\sigma_n} | M] \leq \bar{\lambda}^n / (\bar{\lambda} + u)^n$, which immediately implies that

$$\mathbb{E}^{\vec{\pi}} [e^{-u\sigma_n}] \leq \left(\frac{\bar{\lambda}}{\bar{\lambda} + u} \right)^n.$$

Also, since $1/\sigma_n = \int_0^{\infty} e^{-\sigma_n u} du$, an application of Fubini's theorem together with the last inequality yield

$$(4.7) \quad \mathbb{E}^{\vec{\pi}} \left[\frac{1}{\sigma_n} \right] \leq \int_0^{\infty} \left(\frac{\bar{\lambda}}{\bar{\lambda} + u} \right)^n du = \frac{\bar{\lambda}}{n-1}, \quad n \geq 2.$$

The uniform convergence of k_n to U_0 now follows from (4.5) and (4.6).

Step 2. Here, we will show that $(k_n)_{n \geq 0}$ is a sequence of convex functions. This result would follow from an induction argument once we show that the operator I maps a convex function to a convex function.

Let us assume that $\vec{\pi} \rightarrow w(T, \vec{\pi}, a)$ is a convex function for all $T \geq 0$. Therefore, we can write this convex mapping as $\vec{\pi} \rightarrow w(T-u, \vec{\pi}, a) = \sup_{k \in K_u} \alpha_{k,0}(T-u) + \alpha_{k,1}(T-u)\pi_1 + \dots + \alpha_{k,m}(T-u)\pi_m$, for some constants $\alpha_{k,j} \in \mathbb{R}$ and countable sets K_u . Then using $x_i(t, \vec{\pi}) = m_i(t, \vec{\pi}) / \sum_{j \in E} m_j(t, \vec{\pi})$

and the second equality in (4.1) we obtain

$$\begin{aligned}
(4.8) \quad Iw(T, \vec{\pi}, a) &= \int_0^T e^{-\rho u} \sum_{i \in E} c_i m_i(u, \vec{\pi}) du + \int_0^T e^{-\rho u} \sum_{i \in E} \lambda_i m_i(u, \vec{\pi}) \cdot \\
&\quad \cdot \left[\int_{\mathbb{R}^d} \sup_{k \in K_u} \left(\alpha_{k,0}(T-u) + \sum_{j \in E} \alpha_{k,j}(T-u) \frac{\lambda_j f_j(y) m_j(u, \vec{\pi})}{\sum_{l \in E} \lambda_l f_l(y) m_l(u, \vec{\pi})} \right) f_i(y) \nu(dy) \right] du \\
&= \int_0^T e^{-\rho u} \sum_{i \in E} c_i m_i(u, \vec{\pi}) du \\
&\quad + \int_0^T e^{-\rho u} \left[\int_{\mathbb{R}^d} \sup_{k \in K_u} \left(\sum_{j \in E} [\alpha_{k,j}(T-u) + \alpha_{k,0}(T-u)] \lambda_j f_j(y) m_j(u, \vec{\pi}) \right) \nu(dy) \right] du.
\end{aligned}$$

Since $\vec{\pi} \rightarrow m(u, \vec{\pi})$ is linear in $\vec{\pi}$ (see (2.15)) and the supremum of linear functions is convex, the convexity of $\vec{\pi} \rightarrow Iw(T, \vec{\pi}, a)$ follows. \square

Lemma 4.2. $U_0(T, \vec{\pi}, a)$ is continuous as a function of its first two variables.

Proof. The proof will be carried out in two parts. In the first part we will show that $\vec{\pi} \rightarrow U_0(T, \vec{\pi}, a)$, is Lipschitz on D . In the second part we will show that $T \rightarrow U_0(T, \vec{\pi}, a)$ is Lipschitz uniformly in $\vec{\pi}$. But these two imply that $(T, \vec{\pi}) \rightarrow U_0(T, \vec{\pi}, a)$ is continuous for all $a \in \mathcal{A}$ since

$$\begin{aligned}
(4.9) \quad |U_0(T, \vec{\pi}, a) - U_0(S, \vec{p}, a)| &= |U_0(T, \vec{\pi}, a) - U_0(T, \vec{p}, a) + U_0(T, \vec{p}, a) - U_0(S, \vec{p}, a)| \\
&\leq R(T, a) |\vec{\pi} - \vec{p}| + \tilde{R}(a) |T - S|, \quad \vec{\pi}, \vec{p} \in D; T, S \in \mathbb{R}_+,
\end{aligned}$$

in which $R(T, a)$ and $\tilde{R}(a)$ are the Lipschitz constants above.

Step 1. The idea is to use the convexity of U_0 . Unfortunately, the convexity of $\vec{\pi} \rightarrow U_0(T, \vec{\pi}, a)$ implies that this function is Lipschitz only in the *interior* of D . In what follows we will show that $\vec{\pi} \rightarrow U_0(T, \vec{\pi}, a)$ is the restriction of a convex function $\vec{\pi} \rightarrow \tilde{U}_0(\vec{\pi})$ whose domain is strictly larger than D , which implies the Lipschitz continuity of $\vec{\pi} \rightarrow U_0(T, \vec{\pi}, a)$ also on the boundary of the region D . To this end let us define the functional operator \tilde{I} through its action on a test function w as

$$\tilde{I}w(T, \vec{p}, a) = \int_0^T e^{-\rho u} \sum_{i \in E} m_i(u, \vec{p}) \cdot [C(\vec{x}(u, \vec{p}), a) + \lambda_i \cdot S_i w(T-u, \vec{x}(u, \vec{p}), a)] du,$$

for $\vec{p} \in \tilde{D}, T \in \mathbb{R}_+, a \in \mathcal{A}$ in which

$$\tilde{D} = \left\{ \vec{p} \in \mathbb{R}_+^m : \sum_{i \in E} p_i \leq 2 \right\}.$$

Note that \tilde{I} is nothing but an extension of the operator I we defined in the proof of Lemma 4.1. Let us define

$$\tilde{k}_{n+1}(T, \vec{p}, a) = \tilde{I}\tilde{k}_n(T, \vec{p}, a), \quad \tilde{k}_0(T, \vec{p}, a) = 0, \quad T \in \mathbb{R}_+, \vec{p} \in \tilde{D}, a \in \mathcal{A}.$$

Using the very same arguments as in the proof of Lemma 4.1, we can show that $\vec{p} \rightarrow \tilde{k}_n(T, \vec{p}, a)$ is convex for all n , and this sequence of functions uniformly converges to a convex limit $\vec{p} \rightarrow \tilde{U}_0(T, \vec{p}, a)$. Clearly, $\tilde{k}_n(T, \vec{p}, a) = k_n(T, \vec{p}, a)$ when $\vec{p} \in D$. As a result $U_0(T, \vec{p}, a) = \tilde{U}_0(T, \vec{p}, a)$ on D . Since $\tilde{U}_0(T, \vec{p}, a)$ is locally Lipschitz in the interior of \tilde{D} (as a result of its convexity), we see that $\vec{p} \rightarrow U_0(T, \vec{p}, a)$ is Lipschitz on the compact domain D .

Step 2. The Lipschitz property of U_0 with respect to time (uniformly in $\vec{\pi}$) follows from

$$(4.10) \quad |U_0(T, \vec{\pi}, a) - U_0(S, \vec{\pi}, a)| \leq \mathbb{E}^{\vec{\pi}, a} \left[\int_S^T e^{-\rho t} |C(\vec{\Pi}_t, a)| dt \right] \leq c|T - S|. \quad \square$$

Lemma 4.3. *For all $a \in \mathcal{A}$, $T \in \mathbb{R}_+$, $(W_n(T, \vec{\pi}, \cdot))_{n \in \mathbb{N}}$, defined in (3.20), form a sequence of convex functions. Moreover, for each $a \in \mathcal{A}$ and $n \in \mathbb{N}$, the function $(T, \vec{\pi}) \rightarrow W_n(T, \vec{\pi}, a)$ is continuous.*

Proof. The proof of the convexity of $\vec{\pi} \rightarrow W_n(T, \vec{\pi}, a)$ is similar to the proof of convexity of $\vec{\pi} \rightarrow k_n(T, \vec{\pi}, a)$, which is defined in the proof of Lemma 4.1, see Part II of that proof.

The continuity proof on the other hand parallels the continuity proof for $(T, \vec{\pi}) \rightarrow U_0(T, \vec{\pi}, a)$ which we carried out above. The proof of the uniform Lipschitz continuity of W_n with respect to time is similar to the corresponding proof for U in Lemma 4.5 below. \square

Remark 4.1. The value function U is convex in $\vec{\pi}$, since as a function of $\vec{\pi}$, U is the upper envelope of convex functions (W_n) .

Lemma 4.4. *The value function U is Lipschitz continuous in $\vec{\pi}$,*

$$(4.11) \quad |U(T, \vec{\pi}_1, a) - U(T, \vec{\pi}_2, a)| \leq R(T, a)|\vec{\pi}_1 - \vec{\pi}_2|, \quad \vec{\pi}_1, \vec{\pi}_2 \in D; T \leq T_0; a \in \mathcal{A},$$

where the positive constant R depends on T and a .

Proof. The proof parallels Step 1 of the proof of Lemma 4.2. Again a convex sequence of functions is constructed, converging upwards to an extension of U on \tilde{D} (each element in this sequence is an extension of W_n onto the larger domain.). Here, the convergence is not uniform but monotone. The result still follows since the upper envelope of convex functions is convex, so that the limit is convex and therefore Lipschitz in $\vec{\pi}$ on the original domain D . \square

Lemma 4.5. *The value function U is continuous in T uniformly in the other variables, namely*

$$(4.12) \quad |U(T, \vec{\pi}, a) - U(S, \vec{\pi}, a)| \leq c|T - S|, \quad \text{for any } \vec{\pi} \in D; T, S \in (0, T_0]; a \in \mathcal{A}.$$

Proof. Fix $S > T$. Let ξ^S, ξ^T be ε -optimal strategies for $U(S, \vec{\pi}, a)$ and $U(T, \vec{\pi}, a)$ respectively. Then, taking $\tilde{\xi}^S = \xi^T 1_{[0, T]} + \xi_T^T 1_{(T, S]}$ we have

$$\begin{aligned} U(S, \vec{\pi}, a) - U(T, \vec{\pi}, a) &\geq J^{\tilde{\xi}^S}(S, \vec{\pi}, a) - (J^{\xi^T}(T, \vec{\pi}, a) + \varepsilon) \\ &= \mathbb{E}^{\vec{\pi}, a} \left[\int_T^S e^{-\rho s} C(\vec{\Pi}_s, \xi_s^T) ds \right] - \varepsilon \geq -e^{-\rho T}(S - T)c - \varepsilon. \end{aligned}$$

On the other hand, using the strong Markov property of $(\vec{\Pi}, \xi^S)$,

$$\begin{aligned} U(S, \vec{\pi}, a) - U(T, \vec{\pi}, a) &\leq J^{\xi^S}(S, \vec{\pi}, a) + \varepsilon - J^{\xi^S \cdot 1_{[0, T]}}(T, \vec{\pi}, a) \\ &= \mathbb{E}^{\vec{\pi}, a} \left[\int_T^S e^{-\rho s} C(\vec{\Pi}_s, \xi_s^S) ds - \sum_{k: \tau_k > T} e^{-\rho \tau_k} K(\xi_{k-1}^S, \xi_k^S, \vec{\Pi}_{\tau_k^S}) \right] + \varepsilon \\ &\leq \mathbb{E}^{\vec{\pi}, a} \left[e^{-\rho T} \int_T^S e^{-\rho(s-T)} C(\vec{\Pi}_s, \xi_s^S) ds \right] + \varepsilon \\ &\leq e^{-\rho T}(S - T)c + \varepsilon, \end{aligned}$$

Since ε was arbitrary, we therefore conclude that $|U(T, \vec{\pi}, a) - U(S, \vec{\pi}, a)| \leq c|T - S|$ as desired. \square

Lemma 4.6. *For each $a \in \mathcal{A}$ and n , the function $(T, \vec{\pi}) \rightarrow V_n(T, \vec{\pi}, a)$ is continuous.*

Proof. We proved in Lemma 4.2 that $(T, \vec{\pi}) \rightarrow U_0(T, \vec{\pi}, a)$ is continuous. Furthermore, observe that the operator \mathcal{M} preserves continuity: if for all $a \in \mathcal{A}$, $(T, \vec{\pi}) \rightarrow V(T, \vec{\pi}, a)$ is continuous then for $(T_1, \vec{\pi}_1)$ and $(T_2, \vec{\pi}_2)$ close enough

$$(4.13) \quad |\mathcal{M}V(T_1, \vec{\pi}_1, a) - \mathcal{M}V(T_2, \vec{\pi}_2, a)| \leq \max_{b \in \mathcal{A}, b \neq a} |V(T_1, \vec{\pi}_1, b) - V(T_2, \vec{\pi}_2, b)|$$

is small.

The rest of the proof follows due to the properties of the operator \mathcal{G} in (3.3). Indeed, $\mathcal{G}w(\cdot, \cdot, a)$ defines an optimal stopping problem for $\vec{\Pi}$ with terminal reward function $\mathcal{M}w(\cdot, \cdot, a)$. As shown in Corollary 3.1 of Ludkovski and Sezer [2007] (see also Remark 3.4 in Bayraktar and Sezer [2006]), when $\mathcal{M}w$ is continuous, then the value function $\mathcal{G}w$ of this optimal stopping problem is also continuous. Therefore, by induction, $V_{n+1} = \mathcal{G}V_n$ is continuous. \square

Corollary 4.1. *The value function $U(\cdot, \cdot, a)$ is continuous for all $a \in \mathcal{A}$. Moreover, $(V_n(\cdot, \cdot, a))_{n \in \mathbb{N}}$ defined in (3.4) and $(W_n(\cdot, \cdot, a))_{n \geq 0}$, defined in Proposition 3.4, both converge to $U(\cdot, \cdot, a)$ uniformly for all $a \in \mathcal{A}$.*

Proof. Lemmas 4.4 and 4.5 imply the continuity of $U(\cdot, \cdot, a)$ (also see (4.9)). Now the rest of the statement of the corollary follows from Dini's theorem, which states that pointwise convergence of continuous functions to a continuous limit implies uniform convergence on compacts. \square

Using Corollary 4.1 we obtain the following explicit existence result about an optimal strategy for U :

Proposition 4.1. *Let us extend the value functions U_0 and U so that*

$$(4.14) \quad U_0(T, \vec{\pi}, a) = U(T, \vec{\pi}, a) = 0, \quad T \in [-\varepsilon, 0], \vec{\pi} \in D, a \in \mathcal{A},$$

for some strictly positive constant ε . Let us recursively define a strategy $\xi^* = (\xi_0, \tau_0; \xi_1, \tau_1, \dots)$ via $\xi_0 = a, \tau_0 = 0$ and

$$(4.15) \quad \begin{cases} \tau_{k+1} = \inf\{s \in [\tau_k, T] : U(T-s, \vec{\Pi}(s), \xi_k) = \mathcal{M}U(T-s, \vec{\Pi}(s), \xi_k)\}; \\ \xi_{k+1} = d_{\mathcal{M}U}(T - \tau_{k+1}, \vec{\Pi}(\tau_{k+1}), \xi_k), \quad k = 0, 1, \dots, \end{cases}$$

with the convention that $\inf \emptyset = T + \varepsilon$. Then ξ^* is an optimal strategy for (2.6), i.e.,

$$(4.16) \quad U(T, \vec{\pi}, a) = \mathbb{E}^{\vec{\pi}, a} \left[\int_0^T e^{-\rho s} C(\vec{\Pi}(s), \xi_s^*) ds - \sum_{k: \tau_k \leq T} e^{-\rho \tau_k} K(\xi_k, \xi_{k+1}, \vec{\Pi}(\tau_k)) \right].$$

Proof. We will show that for $n = 1, 2, \dots$

$$(4.17) \quad \begin{aligned} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\tau_n} e^{-\rho s} C(\vec{\Pi}(s), \xi_s) ds - \sum_{k=0}^{n-1} e^{-\rho \tau_k} K(\xi_k, \xi_{k+1}, \vec{\Pi}(\tau_k)) \right] \\ = U(T, \vec{\pi}, a) - \mathbb{E}^{\vec{\pi}, a} \left[e^{-\rho \tau_n} U(T - \tau_n, \vec{\Pi}(\tau_n), \xi_n) \right]. \end{aligned}$$

Suppose that (4.17) is true. Then

$$(4.18) \quad \begin{aligned} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^T e^{-\rho s} C(\vec{\Pi}(s), \xi_s) ds - \sum_{k=0}^{n-1} e^{-\rho \tau_k} K(\xi_k, \xi_{k+1}, \vec{\Pi}(\tau_k)) \right] \\ = U(T, \vec{\pi}, a) - \mathbb{E}^{\vec{\pi}, a} \left[e^{-\rho \tau_n} U(T - \tau_n, \vec{\Pi}(\tau_n), \xi_n) \right] + \mathbb{E}^{\vec{\pi}, a} \left[e^{-\rho \tau_n} U_0(T - \tau_n, \vec{\Pi}(\tau_n), \xi_n) \right]. \end{aligned}$$

Taking the limit as $n \rightarrow \infty$ and using bounded convergence theorem and $\tau_n \rightarrow T + \varepsilon$, we have that

$$\begin{aligned} U(T, \vec{\pi}, a) &= \mathbb{E}^{\vec{\pi}, a} \left[\int_0^T e^{-\rho s} C(\vec{\Pi}(s), \xi_s) ds - \sum_k e^{-\rho \tau_k} K(\xi_k, \xi_{k+1}, \vec{\Pi}(\tau_k)) \right] \\ &\leq \mathbb{E}^{\vec{\pi}, a} \left[\int_0^T e^{-\rho s} C(\vec{\Pi}(s), \xi_s) ds - \sum_{k: \tau_k \leq T} e^{-\rho \tau_k} K(\xi_k, \xi_{k+1}, \vec{\Pi}(\tau_k)) \right], \end{aligned}$$

since $K(a, b, \vec{\pi}) > 0$, and equation (4.16) follows.

To establish (4.17) we proceed by induction. The functions $U(\cdot, \cdot, a)$ and $\mathcal{M}U(\cdot, \cdot, a)$ are continuous by Corollary 4.1. As a result the stopping time

$$(4.19) \quad \tau_1 = \inf \left\{ s \in [0, T] : U(T-s, \vec{\Pi}(s), a) = \mathcal{M}U(T-s, \vec{\Pi}(s), a) \right\},$$

satisfies

$$(4.20) \quad \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\tau_1} e^{-\rho s} C(\vec{\Pi}(s), a) ds + e^{-\rho \tau_1} \mathcal{M}U(T - \tau_1, \vec{\Pi}(\tau_1), a) \right] = U(T, \vec{\pi}, a),$$

see e.g. Proposition 5.12 in Bayraktar et al. [2006]. Rearranging and using $\xi_1 = d_{\mathcal{M}U}(T - \tau_1, \vec{\Pi}(\tau_1), a)$,

$$(4.21) \quad \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\tau_1} e^{-\rho s} C(\vec{\Pi}(s), \xi_0) ds - e^{-\rho \tau_1} K(\xi_0, \xi_1, \vec{\Pi}(\tau_1)) \right] = U(T, \vec{\pi}, a) - \mathbb{E}^{\vec{\pi}, a} \left[e^{-\rho \tau_1} U(T - \tau_1, \vec{\Pi}(\tau_1), \xi_1) \right],$$

proving (4.17) for $n = 1$. Perhaps we should emphasize the dependence on T on the left-hand-side of (4.21) by inserting T as another superscript above \mathbb{E} (we are conditioning on the strong Markov process $t \rightarrow (T - t, \vec{\Pi}_t, \xi_t)$). Although we are not going to implement this for notational consistency/convenience, one should keep this point in mind when reading the rest of the proof.

Assume now that for some $n \geq 1$ (4.17) is satisfied; we will prove that it also holds when we replace n by $n + 1$. Since τ_n 's are all hitting times we have that $\tau_{n+1} = \tau_n + \tau_1 \circ \theta_{\tau_n}$.

$$(4.22) \quad \begin{aligned} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\tau_{n+1}} e^{-\rho s} C(\vec{\Pi}(s), \xi_s) ds - \sum_{k=0}^n e^{-\rho \tau_k} K(\xi_k, \xi_{k+1}, \vec{\Pi}(\tau_k)) \right] &= \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{\tau_n} e^{-\rho s} C(\vec{\Pi}(s), a) ds \right. \\ &\quad \left. - \sum_{k=0}^{n-1} e^{-\rho \tau_k} K(\xi_k, \xi_{k+1}, \vec{\Pi}(\tau_k)) + e^{-\rho \tau_n} \mathbb{E}^{\vec{\Pi}(\tau_n), \xi_n} \left[\int_0^{\tau_1} e^{-\rho s} C(\vec{\Pi}(s), \xi_0) ds - e^{-\rho \tau_1} K(\xi_0, \xi_1, \vec{\Pi}(\tau_1)) \right] \right] \end{aligned}$$

Using (4.21) we can then write

$$(4.23) \quad \begin{aligned} \mathbb{E}^{\vec{\pi}, a} \left[e^{-\rho \tau_n} \mathbb{E}^{\vec{\Pi}(\tau_n), \xi_n} \left[\int_0^{\tau_1} e^{-\rho s} C(\vec{\Pi}(s), \xi_0) ds - e^{-\rho \tau_1} K(\xi_0, \xi_1, \vec{\Pi}(\tau_1)) \right] \right] \\ = \mathbb{E}^{\vec{\pi}, a} \left[e^{-\rho \tau_n} U(T - \tau_n, \vec{\Pi}(\tau_n), \xi_n) - e^{-\rho \tau_{n+1}} U(T - \tau_{n+1}, \vec{\Pi}(\tau_{n+1}), \xi_{n+1}) \right]. \end{aligned}$$

Using (4.22) and (4.23) together with the induction hypothesis, we obtain (4.17) when n is replaced by $n + 1$. \square

Let

$$(4.24) \quad \begin{aligned} \mathcal{C}_s(a) &\triangleq \{ \vec{\pi} \in D : U(s, \vec{\pi}, a) > \mathcal{M}U(s, \vec{\pi}, a) \}, \\ \Gamma_s(a) &\triangleq \{ \vec{\pi} \in D : U(s, \vec{\pi}, a) = \mathcal{M}U(s, \vec{\pi}, a) \} \end{aligned}$$

denote the continuation and switching regions for initial policy a with s time units until maturity. The switching region can further be decomposed as the union $\cup_{b \in \mathcal{A}} \Gamma_s(a, b)$ of the regions defined as

$$(4.25) \quad \Gamma_s(a, b) \triangleq \{ \vec{\pi} \in D : U(s, \vec{\pi}, a) = U(s, \vec{\pi}, b) - K(a, b, \vec{\pi}) \}, \quad b \in \mathcal{A},$$

The results in the previous section imply that to solve (3.13) with initial horizon of T , one maintains the initial policy a and observes the process $\vec{\Pi}$ until time $\tau_1 = \tau_1(T)$, whence it enters the region $\Gamma_{T-\tau_1}(a)$. At this time, if $\vec{\Pi}_{\tau_1}$ is in the set $\Gamma_{T-\tau_1}(a, b)$ we take $\xi_1 = b$; that is, we select the b 'th policy in the policy set \mathcal{A} . The boundaries of $\Gamma_s(a, b)$ are termed switching boundaries and provide an efficient way of summarizing the optimal strategy of the controller. We plot these curves in our examples in Section 6.

5. EXTENSIONS

5.1. Infinite Horizon Formulation. In many practical settings, the controller does not have a natural horizon for her strategies. In such cases it is more appropriate to consider infinite-horizon setting. Due to time-homogeneity, the infinite-horizon problem is stationary in time, reducing the dimension by one. In particular, the optimal strategy can be simplified with a single switching-boundary plot, as $\Gamma_s(a)$'s are independent of s .

For $\rho > 0$, let

$$(5.1) \quad V_\rho(\vec{\pi}, a) = \sup_{\xi \in \mathcal{U}(\infty)} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^\infty e^{-\rho t} C(\vec{\Pi}(t), \xi_t) dt - \sum_k e^{-\rho \tau_k} K(\xi_{k-1}, \xi_k, \vec{\Pi}(\tau_k)) \right].$$

Here $\mathcal{U}(\infty)$ denotes the admissible strategies that satisfy $\mathbb{E}^{\vec{\pi}, a} \left[\sum_k e^{-\rho \tau_k} K(\xi_{k-1}, \xi_k, \vec{\Pi}(\tau_k)) \right] < \infty$.

The next proposition shows that the infinite horizon problem can be uniformly approximated by the finite horizon problems. In fact, the convergence is exponentially fast in the time horizon T .

Proposition 5.1. *There exists a constant R such that*

$$(5.2) \quad |U(T, \vec{\pi}, a) - V_\rho(\vec{\pi}, a)| \leq e^{-\rho T} R.$$

Proof. Let ξ^T be an ε -optimal strategy of $U(T, \vec{\pi}, a)$ and $\tilde{\xi}^T = \xi^T(t)1_{[0,T]} + \xi_T^T 1_{(T,\infty)} \in \mathcal{U}(\infty)$. Then

$$\begin{aligned} V_\rho(\vec{\pi}, a) - U(T, \vec{\pi}, a) &\geq \mathbb{E}^{\vec{\pi}, a} \left[\int_T^\infty e^{-\rho s} C(\vec{\Pi}_s, \tilde{\xi}_s^T) ds \right] - \varepsilon \\ &\geq -e^{-\rho T} \int_0^\infty e^{-\rho s} c ds - \varepsilon \geq -e^{-\rho T} c/\rho - \varepsilon. \end{aligned}$$

On the other hand, using an ε -optimal control ξ^∞ of $V_\rho(\vec{\pi}, a)$,

$$\begin{aligned} V_\rho(\vec{\pi}, a) - U(T, \vec{\pi}, a) &\leq \mathbb{E}^{\vec{\pi}, a} \left[\int_T^\infty e^{-\rho s} C(\vec{\Pi}_s, \xi_s^\infty) ds - \sum_{k: \tau_k > T} e^{-\rho \tau_k} K(\xi_{k-1}^\infty, \xi_k^\infty, \vec{\Pi}_{\tau_k^\infty}) \right] + \varepsilon \\ &\leq \mathbb{E}^{\vec{\pi}, a} \left[e^{-\rho T} \mathbb{E}^{\vec{\Pi}_T, \xi_T^\infty} \left[\int_0^\infty e^{-\rho s} C(\vec{\Pi}_s, \xi_T^\infty) ds \right] \right] + \varepsilon \\ &\leq e^{-\rho T} \tilde{R} + \varepsilon, \end{aligned}$$

for some constant \tilde{R} where the last line used the fact that the inner term, which is the infinite-horizon counterpart of U_0 , is uniformly bounded on the compact domain $D \times \mathcal{A}$. Taking $R = \max(\tilde{R}, c/\rho)$ the proposition follows. \square

The characterization of the value function of the infinite horizon problem, which we give below, follows along same lines as in Section 4.

Proposition 5.2. *V_ρ is the smallest fixed point of the operator $\hat{L}_\rho(V) \triangleq L_\rho(V, \tilde{\mathcal{M}}V)$ where*

$$L_\rho(V, H)(\vec{\pi}, a) = \sup_{t \geq 0} \mathbb{E}^{\vec{\pi}, a} \left[\int_0^{t \wedge \sigma_1} e^{-\rho s} C(\vec{\Pi}_s, a) ds + 1_{\{t < \sigma_1\}} e^{-\rho t} H(\vec{\Pi}_t, a) + 1_{\{t \geq \sigma_1\}} e^{-\rho \sigma_1} V(\vec{\Pi}_{\sigma_1}, a) \right]$$

and

$$\tilde{\mathcal{M}}V(\vec{\pi}, a) = \max_{b \in \mathcal{A}, b \neq a} \{V(\vec{\pi}, b) - K(a, b, \vec{\pi})\}.$$

Note that \hat{L}_ρ is given by

$$(5.3) \quad \hat{L}_\rho w(\vec{\pi}, a) = \sup_{t \geq 0} \left\{ \left(\sum_{i \in E} m_i(t, \vec{\pi}) \right) \cdot e^{-\rho t} \cdot \tilde{\mathcal{M}}w(\vec{x}(t, \vec{\pi}), a) \right. \\ \left. + \int_0^t e^{-\rho u} \sum_{i \in E} m_i(u, \vec{\pi}) \left[C(\vec{x}(u, \vec{\pi}), a) + \lambda_i \tilde{S}_i w(\vec{x}(u, \vec{\pi}), a) \right] du \right\},$$

where

$$\tilde{S}_i w(\vec{\pi}, a) \triangleq \int_{\mathbb{R}^d} w \left(\left(\frac{\lambda_1 f_1(y) \pi_1}{\sum_{j \in E} \lambda_j f_j(y) \pi_j}, \dots, \frac{\lambda_m f_m(y) \pi_m}{\sum_{j \in E} \lambda_j f_j(y) \pi_j} \right), a \right) f_i(y) \nu(dy), \quad i = 1, \dots, m,$$

for a bounded function $w(\cdot, \cdot)$ defined on $D \times \mathcal{A}$ only. The optimal stopping time for V_ρ is now the first entrance time $\tau_0(\vec{\pi})$ of the process $\vec{\Pi}$ to the time-stationary region

$$\Gamma(a) = \left\{ \vec{\pi} \in D : V_\rho(\vec{\pi}, a) = \tilde{\mathcal{M}}V_\rho(\vec{\pi}, a) \right\}.$$

To compute V_ρ we define again

$$W_0(\vec{\pi}, a) = \mathbb{E}^{\vec{\pi}, a} \left[\int_0^\infty e^{-\rho s} C(\vec{\Pi}_s, a) ds \right], \quad \text{and} \quad W_{n+1} = \hat{L}_\rho W_n.$$

Then as in Section 4, it can be shown that $W_n \nearrow V_\rho$, and W_n can be computed numerically by using (5.3).

5.2. Costs Incurred at Arrival Times. In many practical settings the arrivals of X are themselves costly which leads us to consider a running cost structure of the form

$$\sum_{j=1}^{N(t)} e^{-\rho\sigma_j} c_i(Y_j, a) 1_{\{M_{\sigma_j}=i\}},$$

where $c_i : \mathbb{R}^d \times \mathcal{A} \mapsto \mathbb{R}$ (with $\int_{\mathbb{R}^d} c_i^+(y, a) \nu_i(dy) < \infty$ for all $i \in E, a \in \mathcal{A}$) is the cost incurred upon an arrival of size Y_j when the controller has policy a in place and the environment is $M_{\sigma_j} = i$. Above $N(t)$ is the number of arrivals by time t , and (σ_j, Y_j) are the arrival times and marks respectively. As an example, see Section 6.3 below.

In the latter case, setting $C(y, \vec{\pi}, a) = \sum_i \pi_i c_i(y, a)$ one deals with the objective function

$$(5.4) \quad \tilde{U}(T, \vec{\pi}, a) \triangleq \sup_{\xi \in \mathcal{U}(T)} \mathbb{E}^{\vec{\pi}, a} \left[\sum_{j=1}^{N(T)} e^{-\rho\sigma_j} C(Y_j, \vec{\Pi}(\sigma_j), \xi_{\sigma_j}) - \sum_k e^{-\rho\tau_k} K(\xi_{k-1}, \xi_k, \vec{\Pi}(\tau_k)) \right],$$

by solving the equivalent coupled stopping problem

$$\tilde{U}(T, \vec{\pi}, a) \triangleq \sup_{\tau \in \mathcal{S}(T)} \mathbb{E}^{\vec{\pi}, a} \left[\sum_{j=1}^{N(\tau)} e^{-\rho\sigma_j} C(Y_j, \vec{\Pi}(\sigma_j), a) + e^{-\rho\tau} \mathcal{M}\hat{U}(T - \tau, \vec{\Pi}(\tau), a) \right],$$

as in Proposition 2.6. One can easily verify that the function \tilde{U} is the smallest fixed point greater than U_0 of the operator \tilde{L} whose action on a test function w is

$$\begin{aligned} \tilde{L}w(T, \vec{\pi}, a) = & \sup_{0 \leq t \leq T} \left\{ \left(\sum_{i \in E} m_i(t, \vec{\pi}) \right) \cdot e^{-\rho t} \cdot \mathcal{M}w(T - t, \vec{x}(t, \vec{\pi}), a) \right. \\ & \left. + \int_0^t e^{-\rho u} \sum_{i \in E} m_i(u, \vec{\pi}) \cdot \lambda_i \left(\int_{\mathbb{R}^d} C(y, \vec{x}(u, \vec{\pi}), a) \nu_i(dy) + S_i w(T - u, \vec{x}(u, \vec{\pi}), a) \right) du \right\}. \end{aligned}$$

6. NUMERICAL ILLUSTRATIONS

Below we provide numerical examples illustrating our model based on the applications outlined in Section 1.1. The numerical implementation proceeds by discretizing the time horizon $[0, T]$ and then directly finding the deterministic supremum over t 's in (3.23). Similarly, the domain D is also discretized and linear interpolation is used for evaluating the jump operator S of (3.24). Because the algorithm proceeds forward in time with $t = 0, \Delta t, \dots, T$, for a given time-step $t = m\Delta t$, the right-hand-side in (3.23) is known and one may obtain $U(m\Delta t, \vec{\pi}, a)$ directly.

On infinite horizon since there is no time-variable the dynamic programming equation (5.3) is coupled. Accordingly, one must use the iterative sequence of W_n , as detailed in Section 5.1. Namely, one first computes $W_0 = U_0$ by iterating (4.2), and then applies \hat{L}_ρ several times to find a suitably good approximation W_n .

6.1. Optimal Tracking of ‘On-Off’ System. Consider a physical system (for example a military radar) that can be in two states $E = \{1, 2\}$. Information about the system is obtained via a point process X that summarizes observations. The controller wishes to track the state of the system by announcing at each time $0 \leq t \leq T$ whether the current state is $a = 1$ or $a = 2$, $\mathcal{A} = \{1, 2\}$. The controller faces a penalty if her announcement is incorrect; namely a running benefit is assessed at rate $c_1(1) dt$ (respectively, $c_2(2) dt$) if the controller declares $\xi_t = 1$ and indeed $M_t = 1$ (resp. $\xi_t = 2$ and $M_t = 2$). If the controller is incorrect then no benefit is received. Moreover, the controller faces fixed costs $K(1, 2)$ (resp. $K(2, 1)$) from switching her announcement from state 1 to state 2. $K(a, b)$ ’s represent the effort for disseminating new information, alerting other systems, triggering event protocols, etc. A case in point is the alert announcements by the Department of Homeland Security regarding terrorist threat level which receive major coverage in the media and have significant nationwide implications with high associated costs. Thus, both in the case of an upgrade and in the case of a downgrade, specific protocols must be followed by appropriate government and corporate departments. These effects imply that alert levels should be changed only when significant changes occur in the controller beliefs.

To illustrate we take without loss of generality $c_1(1) = c_2(2) = 1$, $c_1(2) = c_2(1) = 0$ and first consider $K(1, 2) = K(2, 1) = 0.05$, $\rho = 0$, $T = 1$. We assume that X is a simple Poisson process with corresponding intensities $\lambda(M) = [1, 4]$, so that arrivals are much more likely in the ‘alarm’ state 2. Finally, the generator of M is

$$Q = \begin{pmatrix} -1 & 1 \\ 3 & -3 \end{pmatrix},$$

so that on average an alarm should be declared $\lim_{t \rightarrow \infty} \mathbb{P}\{M_t = 2\} = 25\%$ of the time.

Figure 1 shows the results, in particular the switching regions $\Gamma_s(a, b)$. We observe a highly non-trivial dependence of the switching boundaries on time to maturity. First, very close to maturity, no switching takes place at all, as the fixed switching costs K dominate any possible gain to be made. For small s , the no-switching region $\mathcal{C}_s(a)$ is very large, because the controller is reluctant to change her announcement close to maturity. On the other hand, we observe that the switching region in policy 1 narrows between medium $s \sim 0.2$ and large s . This happens again due to the finite horizon. With $s = 0.2$, when the controller believes that $M_t = 2$ with high probability, it is unlikely that M_t will change again before maturity, so that the optimal strategy is to pay the switching cost $K(1, 2)$ and plan to maintain policy 2 until expiration. On the other hand, for large $s \geq 0.5$, even when $\mathbb{P}\{M_t = 2\} = 1 - \pi_1$ is quite large, the controller knows that soon enough M is likely to return to state 1 (since $q_{2,1}$ is large); rather than do two switches and track M , the controller takes a shortcut and continues to maintain policy 1 (with the knowledge that her error is likely to be shortlived). This “shortcircuiting” will disappear only when π_1 is extremely

small. Note that this phenomenon is one-sided: because $q_{1,2}$ is small, the upper boundary $\Gamma_s(2, 1)$ is monotonically decreasing over time.




FIGURE 1. Sequential tracking of a two-state Markov chain. The left panel shows the value functions $U(T, \vec{\pi}, \cdot)$, $a \in \{1, 2\}$, as a function of π_1 for $T = 1$. Recall that in this case $D = \{(\pi_1, 1 - \pi_1) : 0 \leq \pi_1 \leq 1\}$. The vertical lines indicate the boundary of $\Gamma_T(1, 2)$ and $\Gamma_T(2, 1)$. The right panel shows the switching regions $\Gamma_s(a, b)$ (namely $\Gamma_s(1, 2)$ is below the lower curve and $\Gamma_s(2, 1)$ is above the higher curve) as a function of time to maturity s .

6.2. Policy Making Example. The Federal Reserve Board (the Fed) has the task of adjusting the US monetary policy in response to economic events. The Fed has authority over the overnight interest rates and attempts to implement a loose monetary policy when the economy is weak, and a tight monetary policy when the economy is overheating. Unfortunately, the current state of the economy M is never precisely known; thus the main task of the Fed is to estimate M from various economic information it collects. When the beliefs of the Fed change sufficiently, it will adjust its monetary policy ξ . Such adjustments are expensive, since they are closely followed by market participants and send out important signals to economic agents. Thus, beyond trying to track M , the Fed also seeks *stability* in its policies, in order not to disrupt planning activities of businesses.

As can be seen from this description, this problem fits well into our tracking paradigm of (1.3). For concreteness, let $M = \{M_t\}_{t \geq 0}$ represent the current economy with state space $E = \{1, 2, 3\} \equiv \{\text{Overheating, Growth, Recession}\}$. The generator of M is taken to be

$$Q = \begin{pmatrix} -4 & 3 & 1 \\ 2 & -4 & 2 \\ 0 & 3 & -3 \end{pmatrix}.$$




FIGURE 2. Value function $U(T, \vec{\pi}, a)$ of the Fed policy-making example in Section 6.2 plotted together with the switching regions $\Gamma_T(a, \cdot)$ for each current policy a .

Thus, M moves randomly between all three states (and we assumed that a recession cannot be immediately followed by overheating). In the face of these three states, the Fed also has three policy levels, namely its action set is $\mathcal{A} = \{0, 1, 2\} = \{\text{Tight}, \text{Normal}, \text{Accommodating}\}$.

The cost function $C(\vec{\pi}, a) = \sum_{i \in E} c_i(a) \pi_i$, is given by the matrix

$$c_i(a) = c_{i,a} = \begin{bmatrix} 2 & -1 & -1 \\ 0 & 2 & 0 \\ -1 & -1 & 0 \end{bmatrix}, \quad i \in E, a \in \mathcal{A}.$$

The switching costs are given by $K(a, b) = 0.05 \cdot 1_{\{a \neq b\}}$ for $a, b \in \mathcal{A}$. The observation process X is a simple Poisson process with M -modulated intensity $\vec{\lambda} = [\lambda_1, \lambda_2, \lambda_3] = [1, 2, 5]$. Thus, the worse the economy state, the more frequent are (negative) events observed by the Fed.

Figure 2 illustrates the obtained results for $T = 4$ and no discounting. The triangular regions in Figure 2 are the state space $D = \{\vec{\pi} \in \mathbb{R}_+^3 : \pi_{Ovr} + \pi_{Gro} + \pi_{Rec} = 1\}$. The respective panels show how the initial switching regions $\Gamma_T(a, \cdot)$ and value functions $U(T, \vec{\pi}, a)$ depend on the current policy a . Observe that because the penalty for not tracking recessions is small, starting out in the ‘Normal’ regime, the Fed will never immediately adopt an ‘Accommodating’ policy, $\Gamma_T(1, 2) = \emptyset$. Similarly, because the penalty for missing an overheating economy is very large, the switching regions into a ‘Tight’ policy are large and conversely, the continuation region $C_T(0)$ is large. Also, observe that the value function appears to be not differentiable at the boundaries. Finally, we stress that because of the final horizon, this problem is again non-time-stationary and the solution (as well as $\Gamma_T(a, \cdot)$) depends on remaining time T .

6.3. Customer Call Center Example. Our last example illustrates the structure of the infinite horizon version together with a different cost structure. We consider a call center application that employs a variable number of servers to answer calls. The calling rate fluctuates and is modulated by the unknown environment variable M . Having more servers decreases the per-call costs, but increases fixed costs related to payroll overhead.

We assume that $M_t \in E = \{\text{Low}, \text{Med}, \text{High}\}$ with a generator

$$Q = \begin{pmatrix} -1 & 1 & 0 \\ 1 & -2 & 1 \\ 0 & 1 & -1 \end{pmatrix}.$$

The observed process X represents the actual received calls and is taken to be a compound Poisson process with intensity $\lambda(M_t)$ and marks Y_1, Y_2, \dots that represent intrinsic call costs. Suppose that $Y \in \{6, 12, 24\}$, and the distribution of Y and λ is M -modulated:

$$\nu_{i,j} = \mathbb{P}\{Y = y_j | M_t = i\} = \begin{pmatrix} 1/4 & 1/2 & 1/4 \\ 1/3 & 1/3 & 1/3 \\ 1/4 & 1/4 & 1/2 \end{pmatrix}; \quad \vec{\lambda} = [1 \ 3 \ 4].$$

Thus, as the manager receives calls, she dynamically updates her beliefs about current state of M based on the intervals between call times and observed call types.

The call center manager can choose one of two strategies, namely she can employ either one or two agents, $a \in \mathcal{A} = \{1, 2\}$. Employing a agents leads to per-call costs of $c_1(Y, a) = -Y/a$ and to continuously-assessed costs of $c_2(Y, a) = -(10 + 20a)$. Thus, when $\mathbb{P}\{M_t = \text{High}\}$ is sufficiently high, it is optimal to employ both agents, otherwise one is sufficient. Finally, switching costs for increasing or decreasing number of agents are set at $K(a, b) = 2$. Note that here all the costs are independent of M (and hence of $\vec{\Pi}$).

We consider an infinite horizon formulation and take $\rho = 0.5$. The parameter ρ measures the trade-off between minimizing immediate costs and having a long-term strategy that takes into account future changes in M . Thus $\rho = 0.5$ means that the horizon of the controller is on the time-scale of two time periods. The overall objective is:

$$\sup_{\xi \in \mathcal{U}(\infty)} \mathbb{E}^{\vec{\pi}, a} \left[\sum_{j=1}^{\infty} e^{-\rho \sigma_j} c_1(Y_j, \xi_{\sigma_j}) + \int_0^{\infty} e^{-\rho t} c_2(\xi_t) dt - \sum_k e^{-\rho \tau_k} K(\xi_{k-1}, \xi_k) \right].$$

Figure 3 shows the results, as well as a computed color-coded sample path of $\vec{\Pi}$ which shows the implemented optimal strategy. The given path has four jumps and three policy changes (two changes occur between jumps when $\vec{\Pi}$ enters $\Gamma(1, 2)$, and one change occurs at an arrival when $\vec{\Pi}$ jumps back into $\Gamma(2, 1)$). Observe that in the absence of new information, $\vec{\Pi}$ converges to the fixed point $\vec{\pi}_{\infty} = [0.7, 0.23, 0.07]$ (the invariant distribution of $e^{Q-\Lambda}$), as can be seen from the flow of the paths in Figure 3.




FIGURE 3. Tracking the regime of a customer call center. We show a sample path of $\vec{\Pi}$ inside the simplex $D = \{(\pi_1, \pi_2, \pi_3) : \pi_i \geq 0, \pi_1 + \pi_2 + \pi_3 = 1\}$, as well as the corresponding optimal strategy. The initial state is $\vec{\Pi}_0 = (0, 1, 0)$ and $\xi_0 = 1$. On this path we have $t \in [0, 4]$ and the arrival pairs (corresponding to jumps of X , recall $\vec{\Pi}$ -dynamics in (2.13)) (σ_ℓ, Y_ℓ) for $\ell = 1, 2, 3, 4$ are $(0.51, 2), (0.66, 3), (1.44, 1), (2.23, 2)$, respectively. The resulting optimal strategy ξ^* is color-coded: dashed line for $\xi_t^* = 1$, solid line for $\xi_t^* = 2$.

REFERENCES

- L. Aggoun. Optimal tracking for an insurance model. *Stochastic Anal. Appl.*, 21(6):1207–1214, 2003. ISSN 0736-2994.
- E. Arjas, P. Haara, and I. Norros. Filtering the histories of a partially observed marked point process. *Stochastic Process. Appl.*, 40(2):225–250, 1992. ISSN 0304-4149.
- E. Bayraktar and S. Dayanik. Poisson disorder problem with exponential penalty for delay. *Mathematics of Operations Research*, 31 (2):217–233, 2006.
- E. Bayraktar and S. Sezer. Quickest detection for a Poisson process with a phase-type change-time distribution. Technical report, University of Michigan, 2006. URL <http://arxiv.org/abs/math/0611563>.
- E. Bayraktar, S. Dayanik, and I. Karatzas. Adaptive Poisson disorder problem. *Annals of Applied Probability*, 16 (3):1190–1261, 2006.
- P. Bremaud. *Point Processes and Queues*. Springer, New York, 1981.
- C. Ceci and A. Gerardi. Partially observed control of a Markov jump process with counting observations: equivalence with the separated problem. *Stochastic Processes and Applications*,

78:245–260, 1998.

- N. Chopin and E. Varini. Particle filtering for continuous-time hidden Markov models. *ESAIM Proceedings*, 19:12–17, 2007. Conférence Oxford sur les méthodes de Monte Carlo séquentielles.
- O. L. V. Costa and M. H. A. Davis. Approximations for optimal stopping of a piecewise-deterministic process. *Math. Control Signals Systems*, 1(2):123–146, 1988. ISSN 0932-4194.
- O. L. V. Costa and M. H. A. Davis. Impulse control of piecewise-deterministic processes. *Math. Control Signals Systems*, 2(3):187–206, 1989. ISSN 0932-4194.
- O. L. V. Costa and C. A. B. Raymundo. Impulse and continuous control of piecewise deterministic Markov processes. *Stochastics Stochastics Rep.*, 70(1-2):75–107, 2000. ISSN 1045-1129.
- J. N. Darroch and K. W. Morris. Passage-time generating functions for continuous-time finite Markov chains. *Journal of Applied Probability*, 5(2):414–426, 1968.
- M. H. A. Davis. *Markov Models and Optimization*. Chapman & Hall, London, 1993.
- R. J. Elliott, L. Aggoun, and J. B. Moore. *Hidden Markov models*, volume 29 of *Applications of Mathematics*. Springer-Verlag, New York, 1995. ISBN 0-387-94364-1. Estimation and control.
- U. Gugerli. Optimal stopping of a piecewise-deterministic Markov process. *Stochastics*, 19(4):221–236, 1986.
- J. Jacod and A. Shiryaev. *Limit Theorems for Stochastic Processes*. Springer-Verlag, Berlin, 1987.
- S. Karlin and H. M. Taylor. *A second course in stochastic processes*. Academic Press, New York, 1981. ISBN 0-12-398650-8.
- S. M. Lenhart and Y. C. Liao. Switching control of piecewise-deterministic processes. *J. Optim. Theory Appl.*, 59(1):99–115, 1988.
- M. Ludkovski and S. Sezer. Finite horizon decision timing with partially observable Poisson processes. Technical report, University of Michigan, 2007.
- G. Mazziotto, L. Stettner, J. Szpirglas, and J. Zabczyk. On impulse control with partial observation. *SIAM J. Control Optim.*, 26(4):964–984, 1988. ISSN 0363-0129.
- P. Muller, B. Sanso, and M. De Iorio. Optimal Bayesian design by inhomogeneous Markov chain simulation. *Journal of the American Statistical Association*, 99:788–798, 2004.
- M. F. Neuts. *Structured Stochastic Matrices of M/G/1 Type and Their Applications*. Marcel Dekker, New York, 1989.
- G. Peskir and A. N. Shiryaev. Solving the Poisson disorder problem. In *Advances in Finance and Stochastics*, pages 295–312. Springer, New York, 2002.
- G. Peskir and A. N. Shiryaev. Sequential testing problems for Poisson processes. *Annals of Statistics*, 28:837–859, 2000.
- A. N. Shiryaev. *Optimal stopping rules*. Springer-Verlag, Berlin, 1978.
- A. G. Tartakovsky, B. L. Rozovskii, R. B. Blažek, and H. Kim. Detection of intrusions in information systems by sequential change-point methods. *Stat. Methodol.*, 3(3):252–293, 2006. ISSN 1572-3127.

(E. Bayraktar) DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MICHIGAN, ANN ARBOR, MI 48109
E-mail address: erhan@umich.edu

(M. Ludkovski) DEPARTMENT OF MATHEMATICS UNIVERSITY OF MICHIGAN, ANN ARBOR, MI 48109
E-mail address: mludkov@umich.edu